Incentive Engineering in the Internet Age

David C. Parkes
Harvard University

Mechanism design theory

- Leonid Hurwicz (1960, 1972)
  - communication system, incentive compatibility

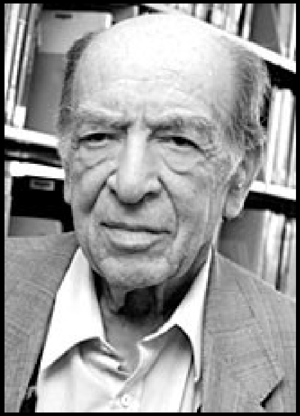

Mechanism design theory

- Leonid Hurwicz (1960, 1972)
  - communication system, incentive compatibility
- Eric Maskin (1977)
  - Nash implementation

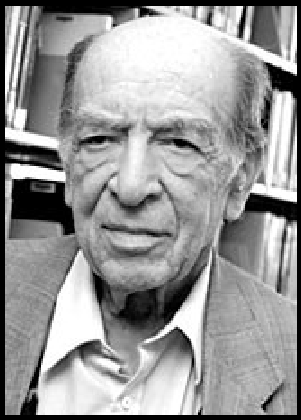

Mechanism design theory

- Leonid Hurwicz (1960, 1972)
  - communication system, incentive compatibility
- Eric Maskin (1977)
  - Nash implementation
- Roger Myerson (1979, 1981)
  - Bayesian mechanism design

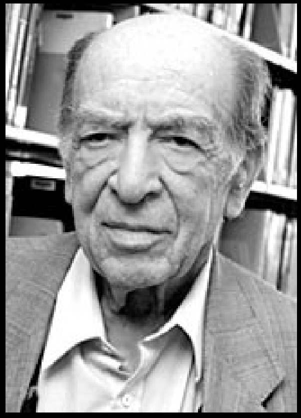

What can be achieved, in principle, by a market system despite agent self interest and private information?

Example: Median Mechanism

(Moulin'80)

Example: Median Mechanism

(Moulin'80)

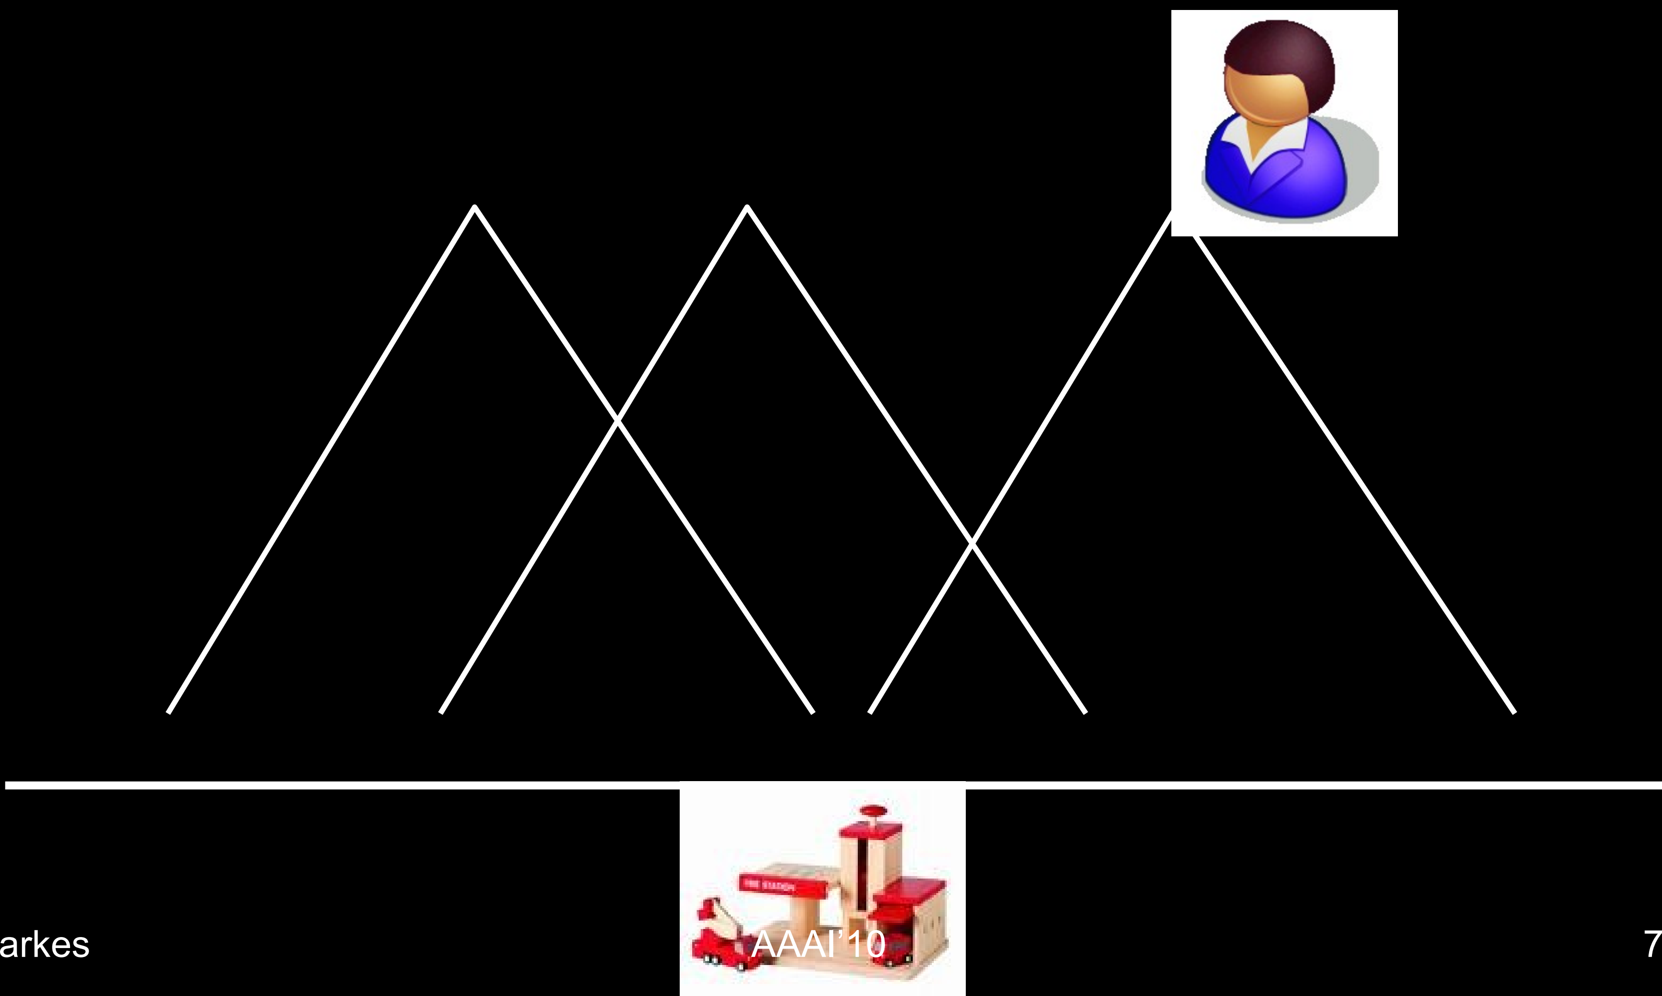

Example: Median Mechanism

(Moulin'80)

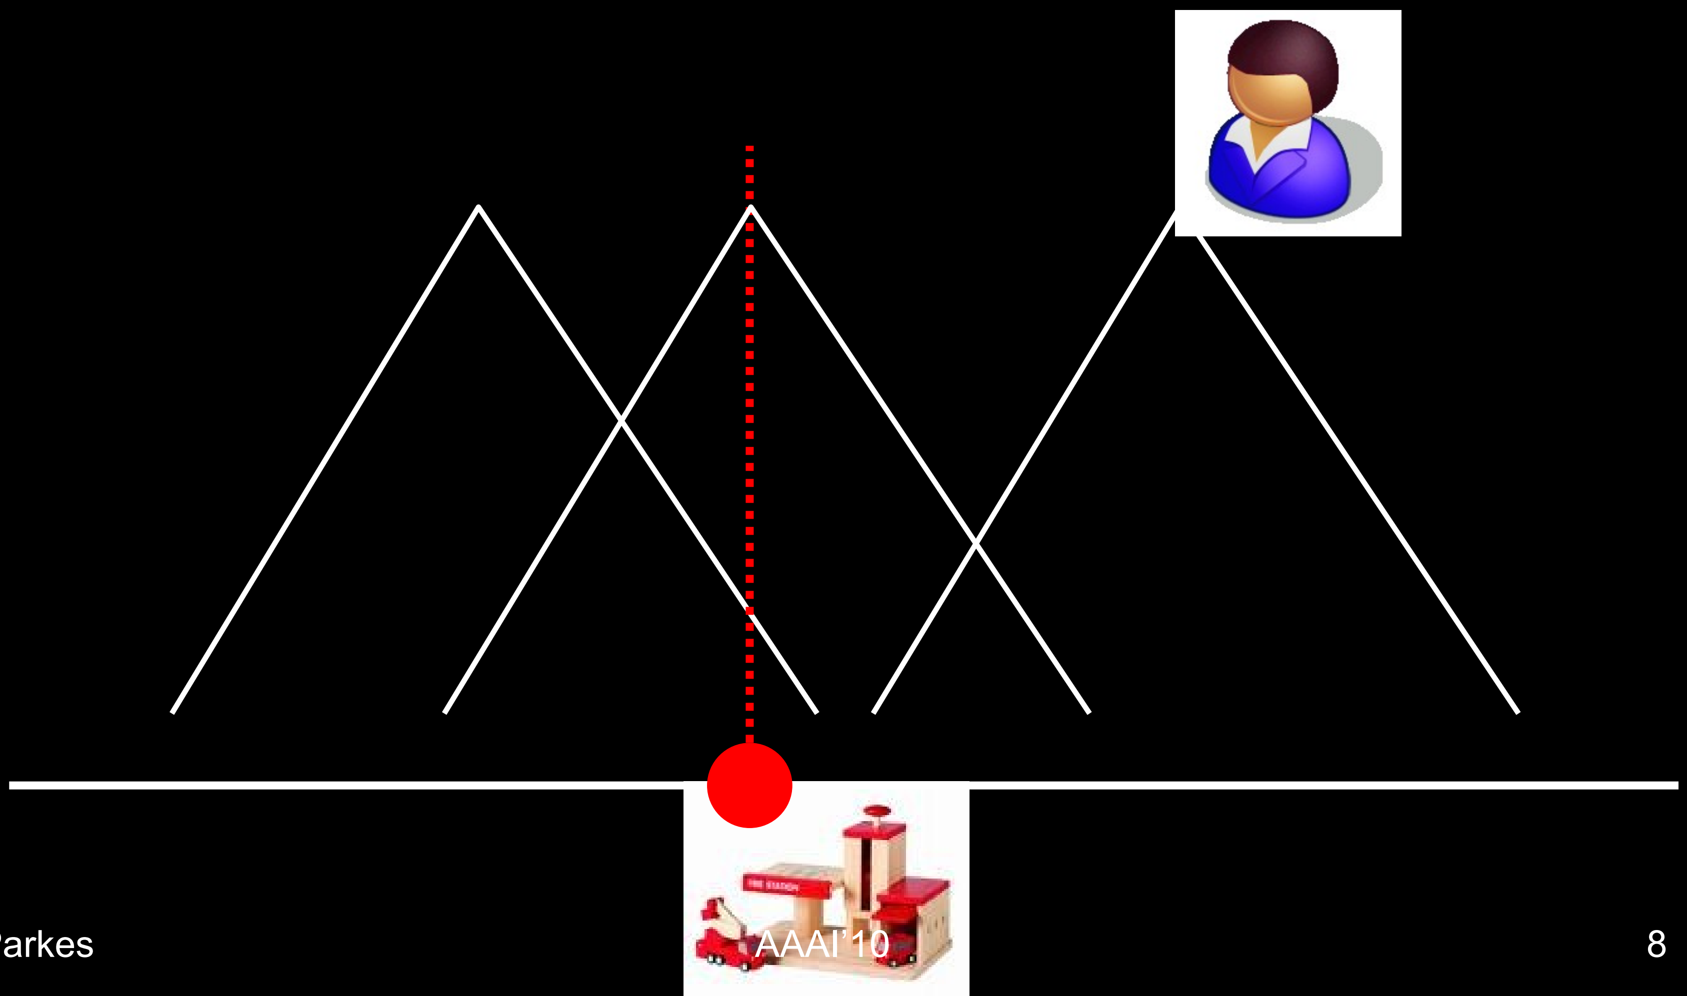

Example: Median Mechanism

(Moulin'80)

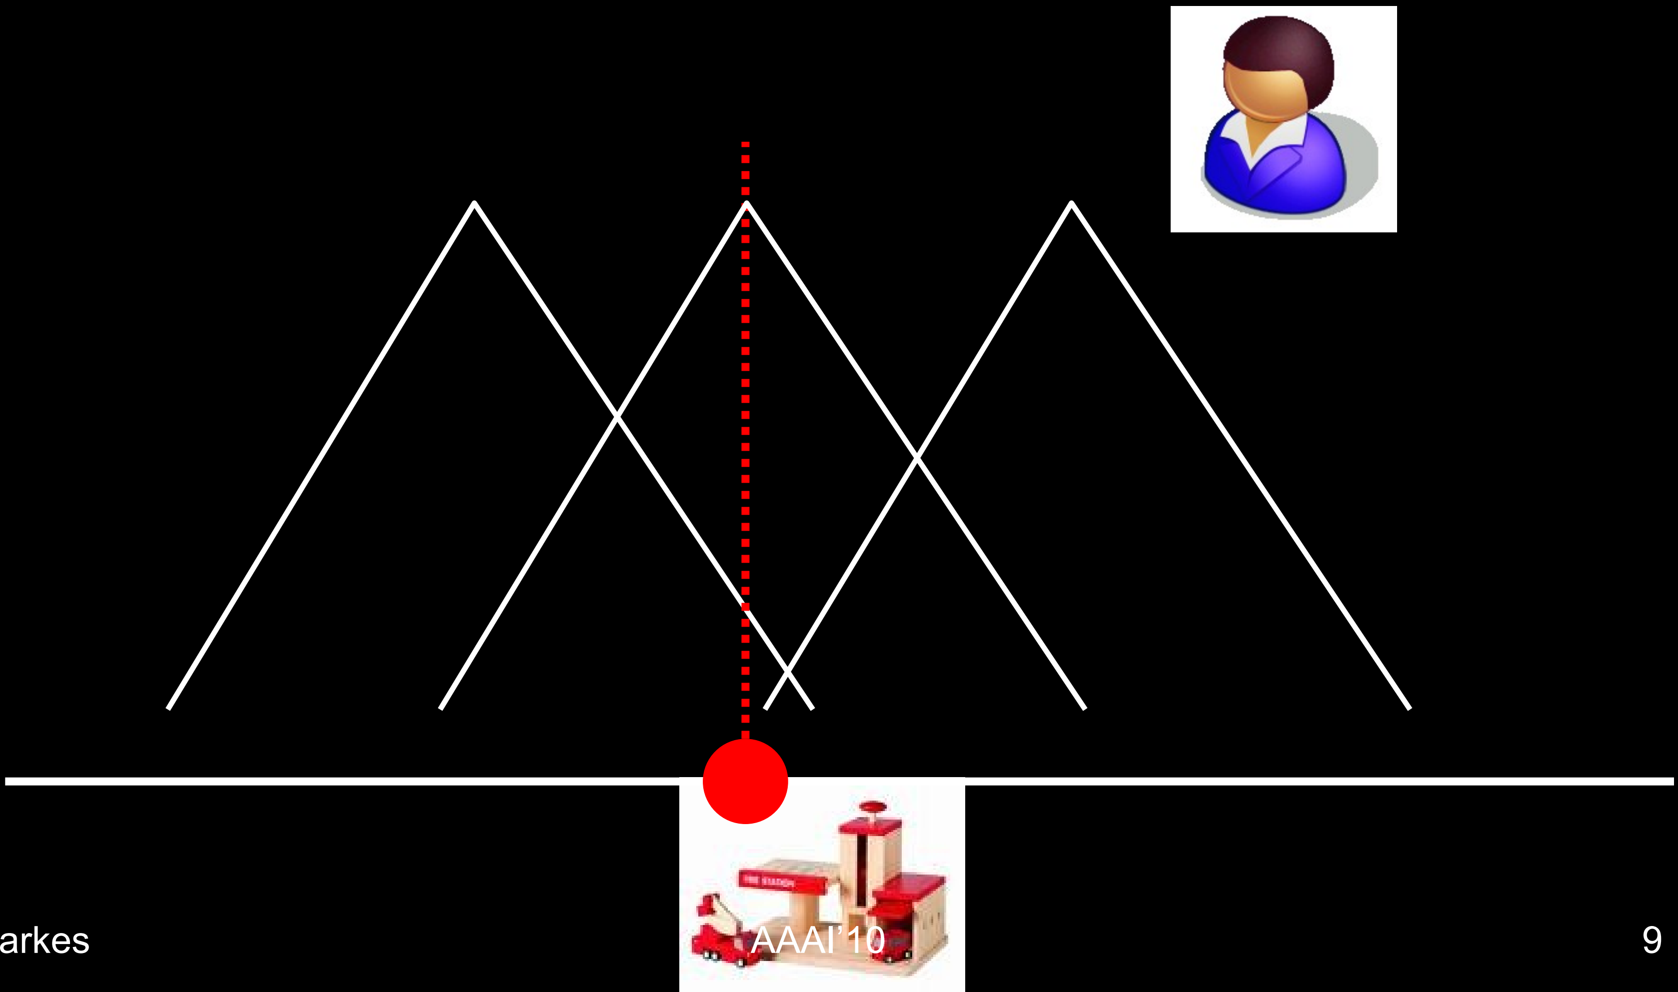

Example: Median Mechanism

(Moulin’80)

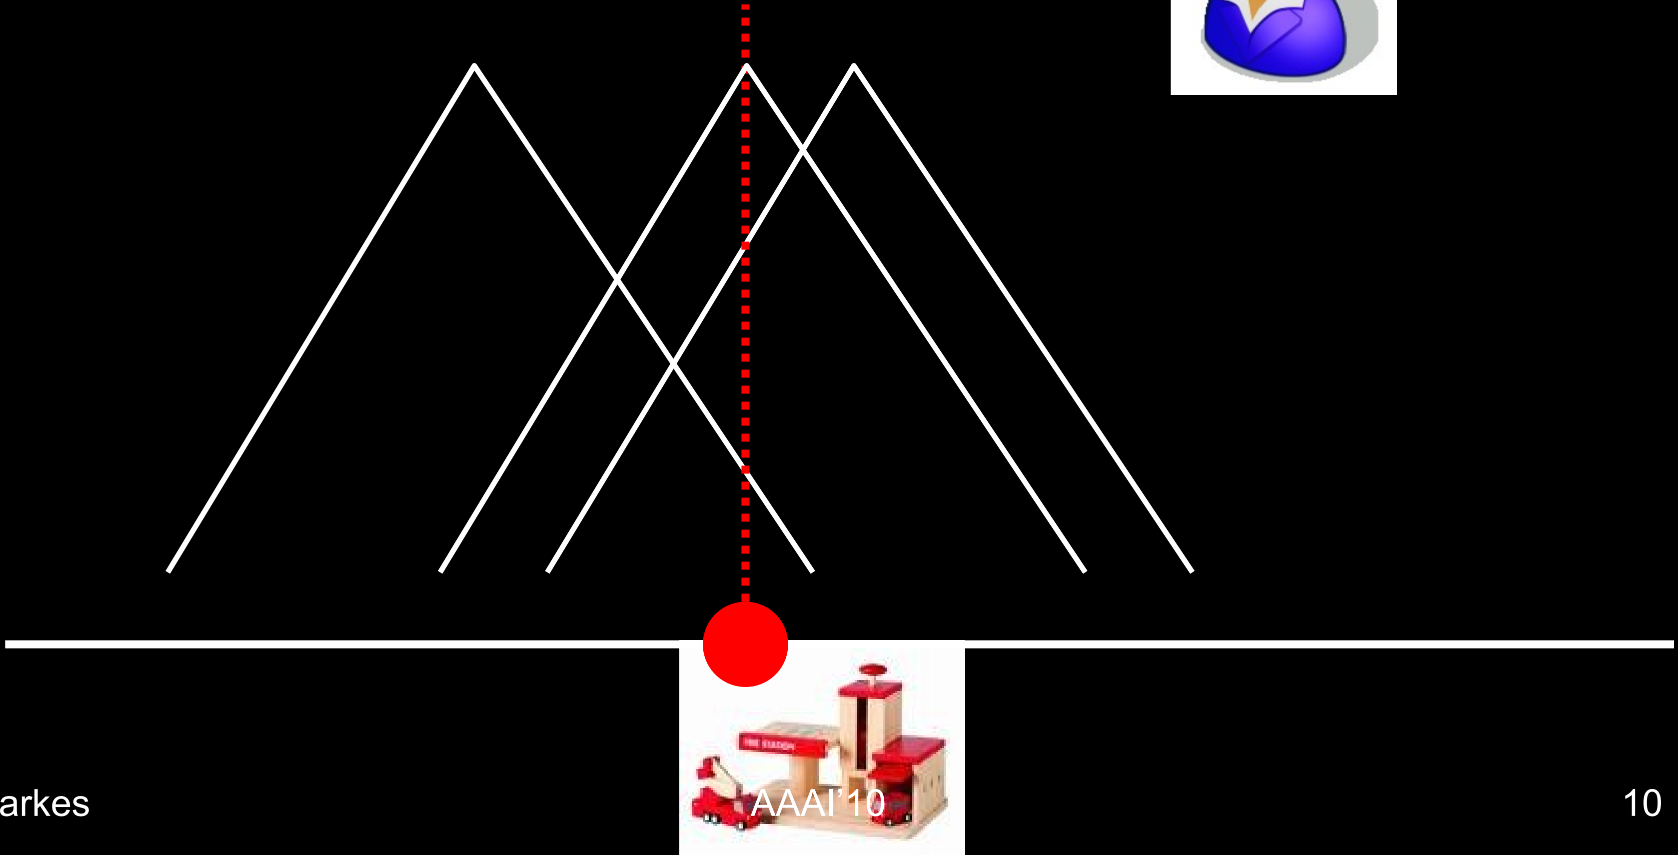

Example: Median Mechanism

(Moulin'80)

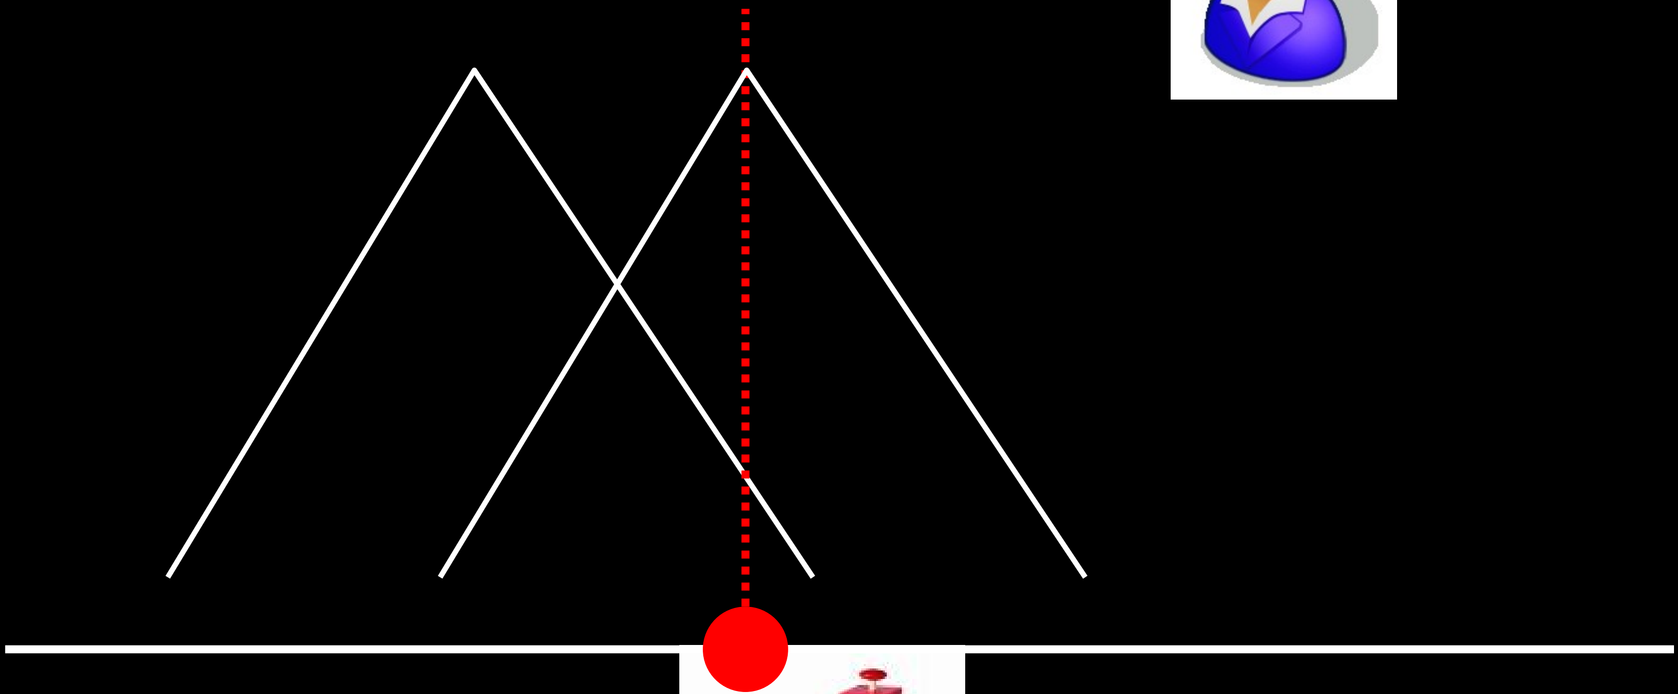

Example: Median Mechanism

(Moulin'80)

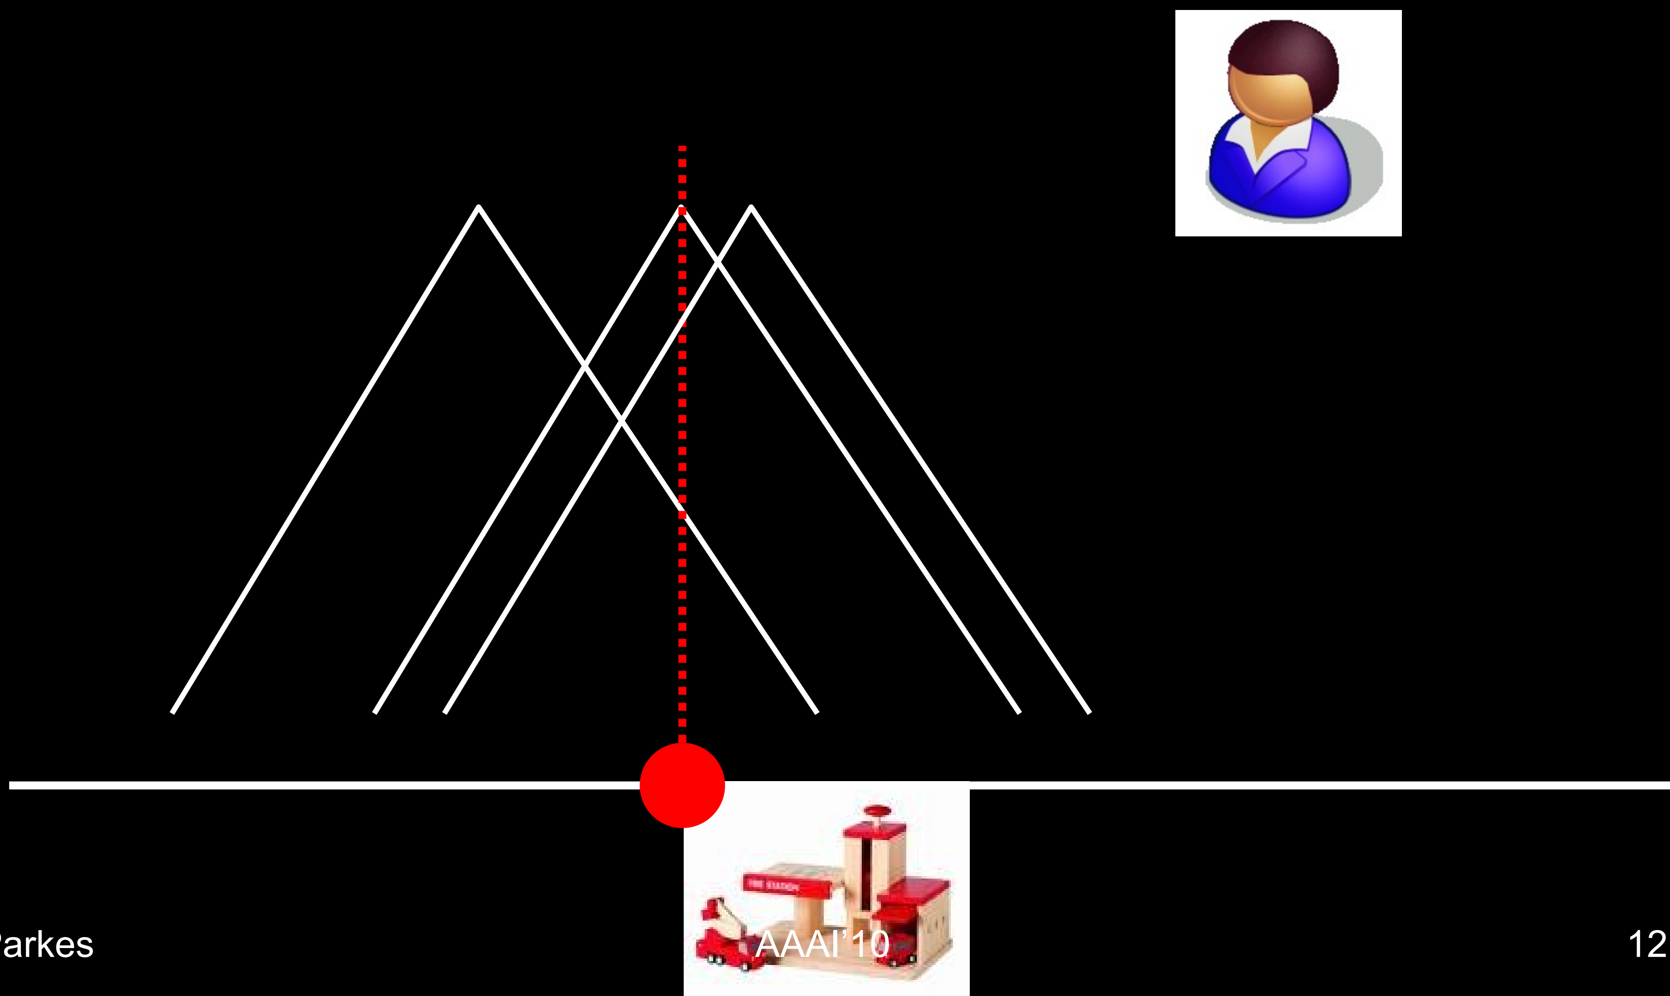

Example: Median Mechanism

(Moulin'80)

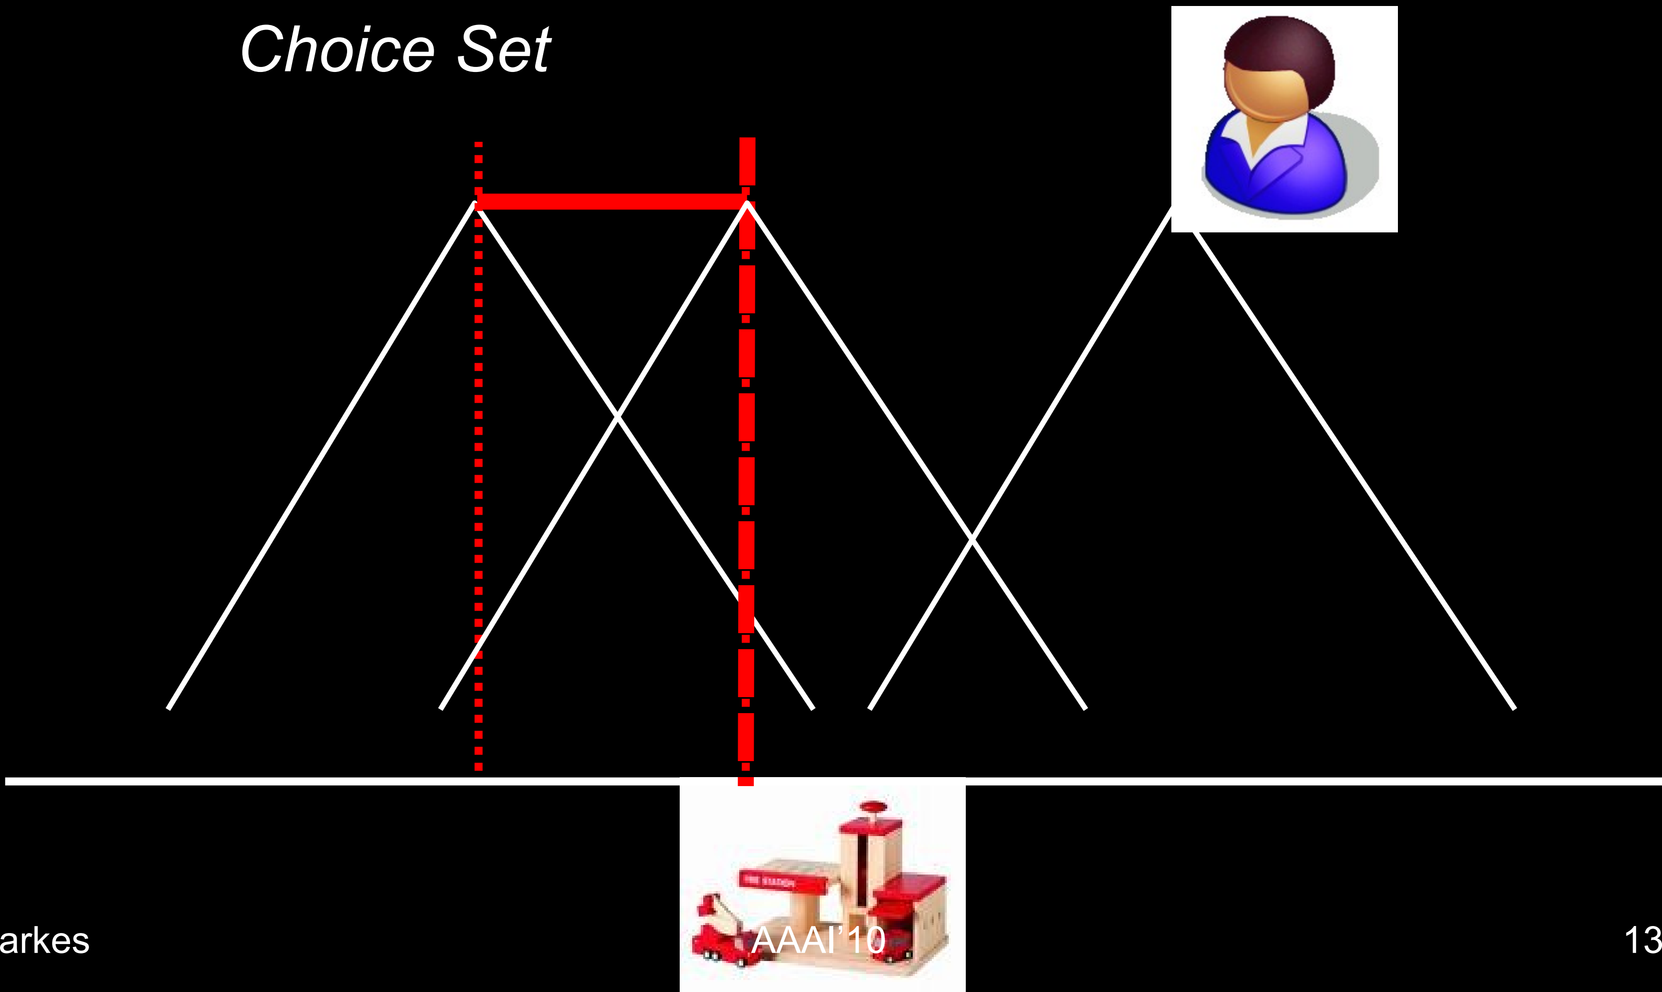

Example: Single item auction

(Vickrey’61)

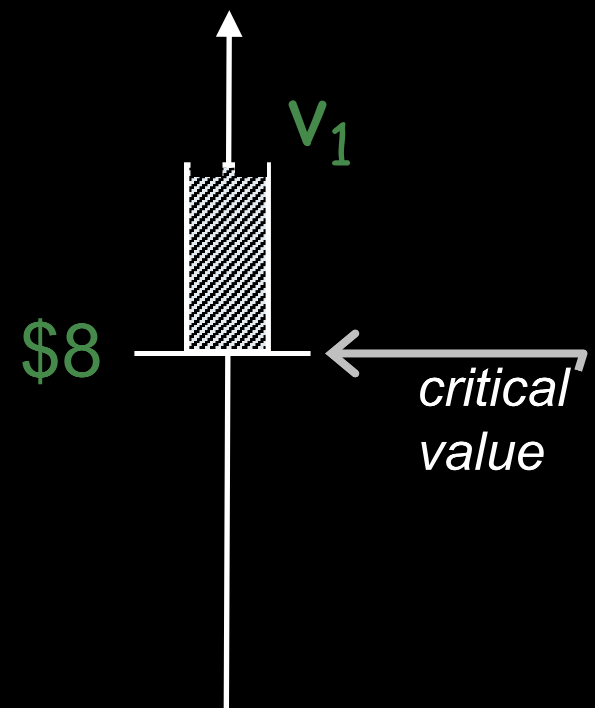

Direct Revelation Mechanism

(Hurwicz'60,'72)

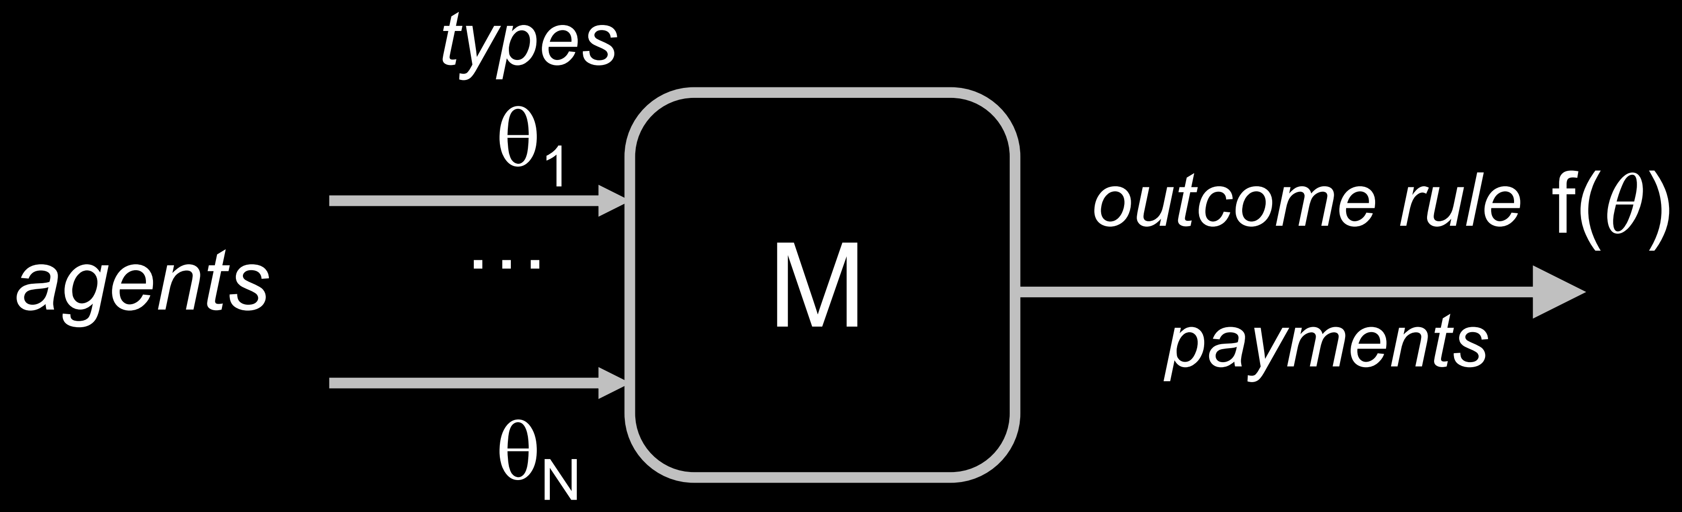

Rules of Encounter

(Rosenschein and Zlotkin 1994; Ephrati and Rosenschein AAAI'91)

- "As distributed systems of computers play an increasingly important role in society, it will be necessary to consider ways in which these machines can be made to interact effectively…

Rules of Encounter

(Rosenschein and Zlotkin 1994; Ephrati and Rosenschein AAAI'91)

- "As distributed systems of computers play an increasingly important role in society, it will be necessary to consider ways in which these machines can be made to interact effectively… *Adjusting the rules of public behavior (the rules of the game) by which the programs must interact can influence the private strategies that designers set up in their machines*."

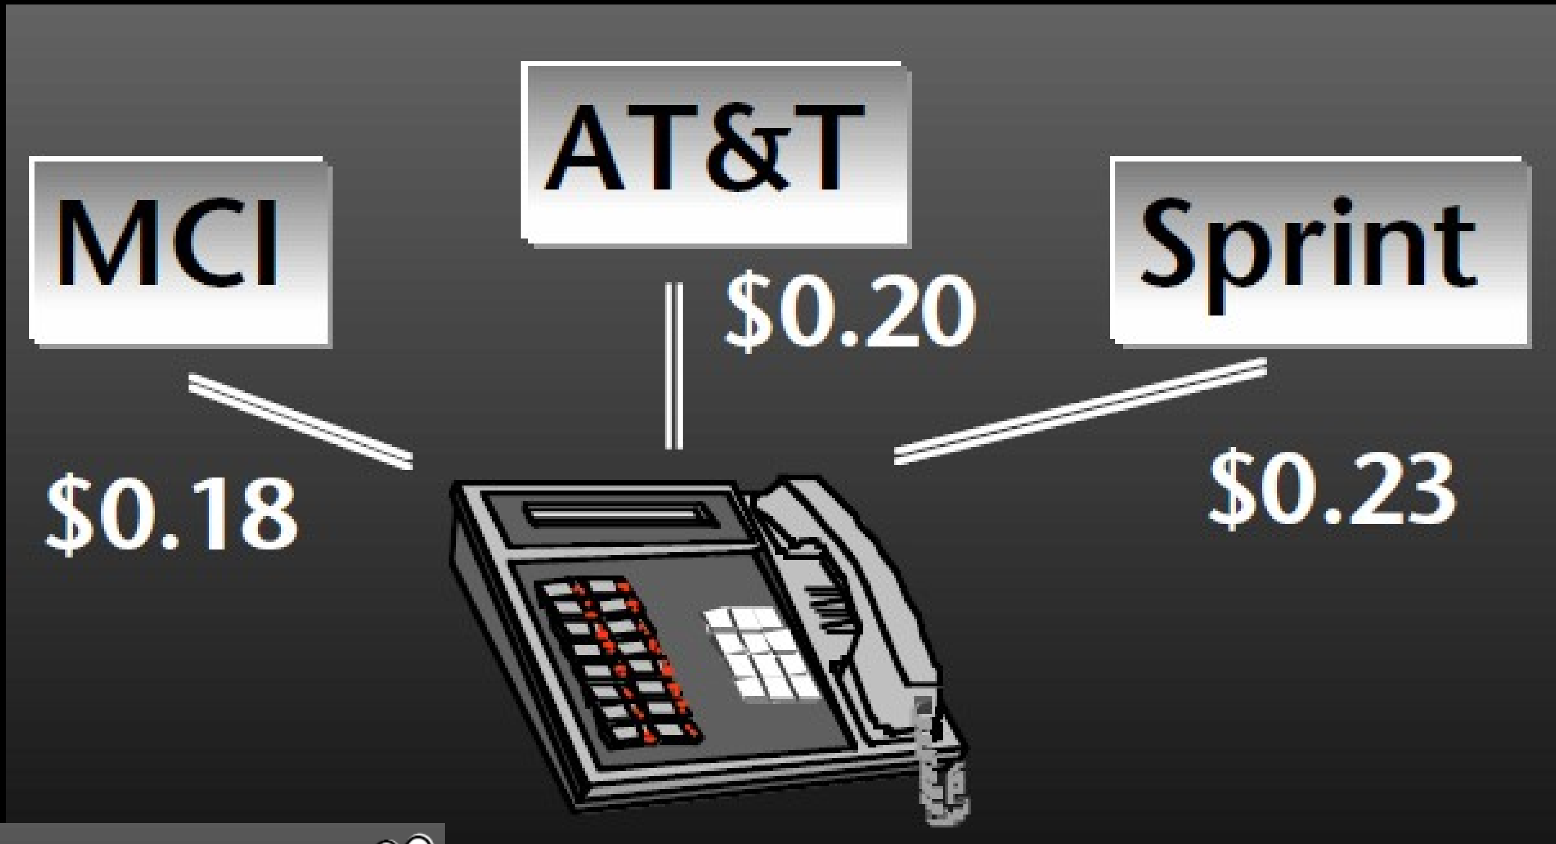
MCI
AT&T
Sprint
$0.20
$0.18
$0.23

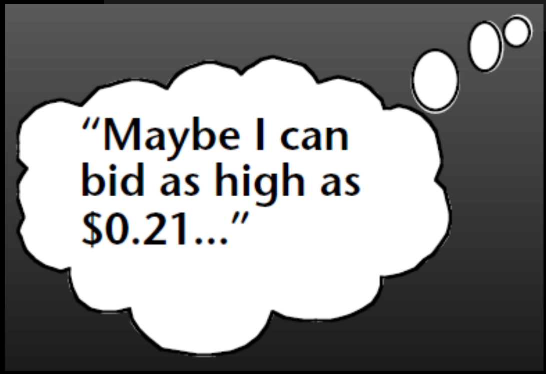
"Maybe I can bid as high as $0.21…"

- “… they’ll pay programmers to develop sophisticated models of their opponents’ bidding strategies… put energy into trying to discover relevant information about their opponents…

  Ultimately, this sort of effort drains resources that might be better spent elsewhere…”

Task negotiation

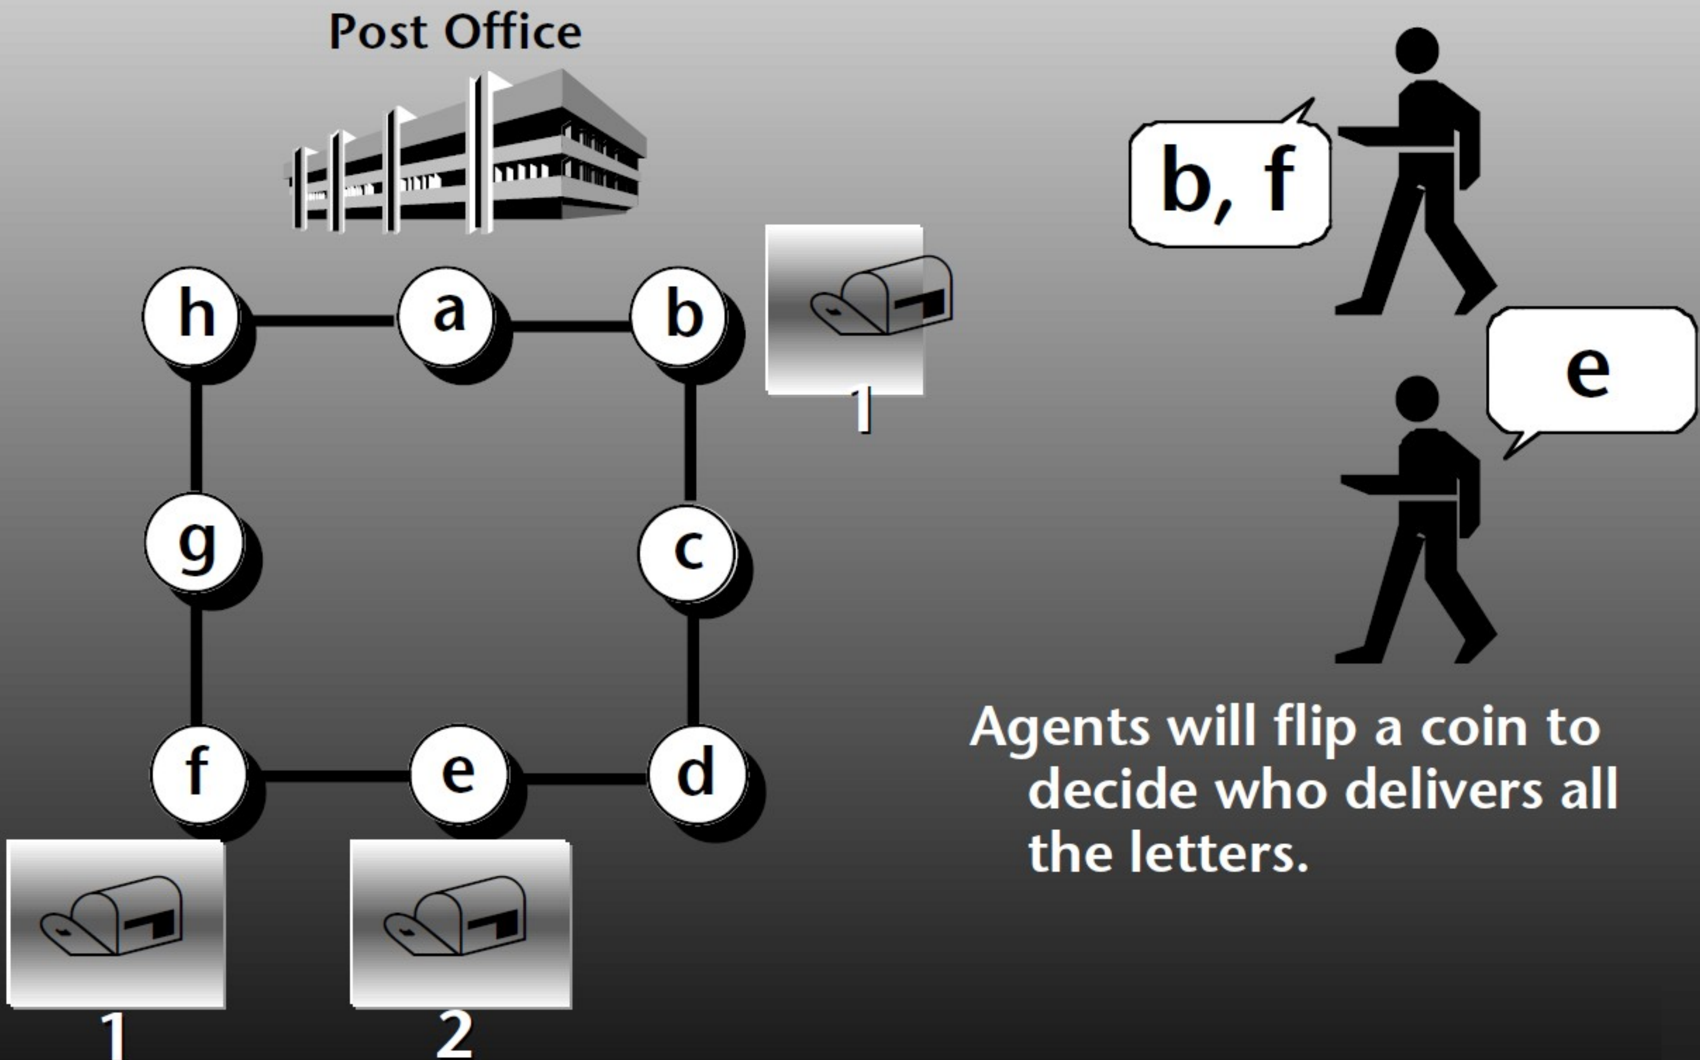

Agents will flip a coin to decide who delivers all the letters.

Task negotiation

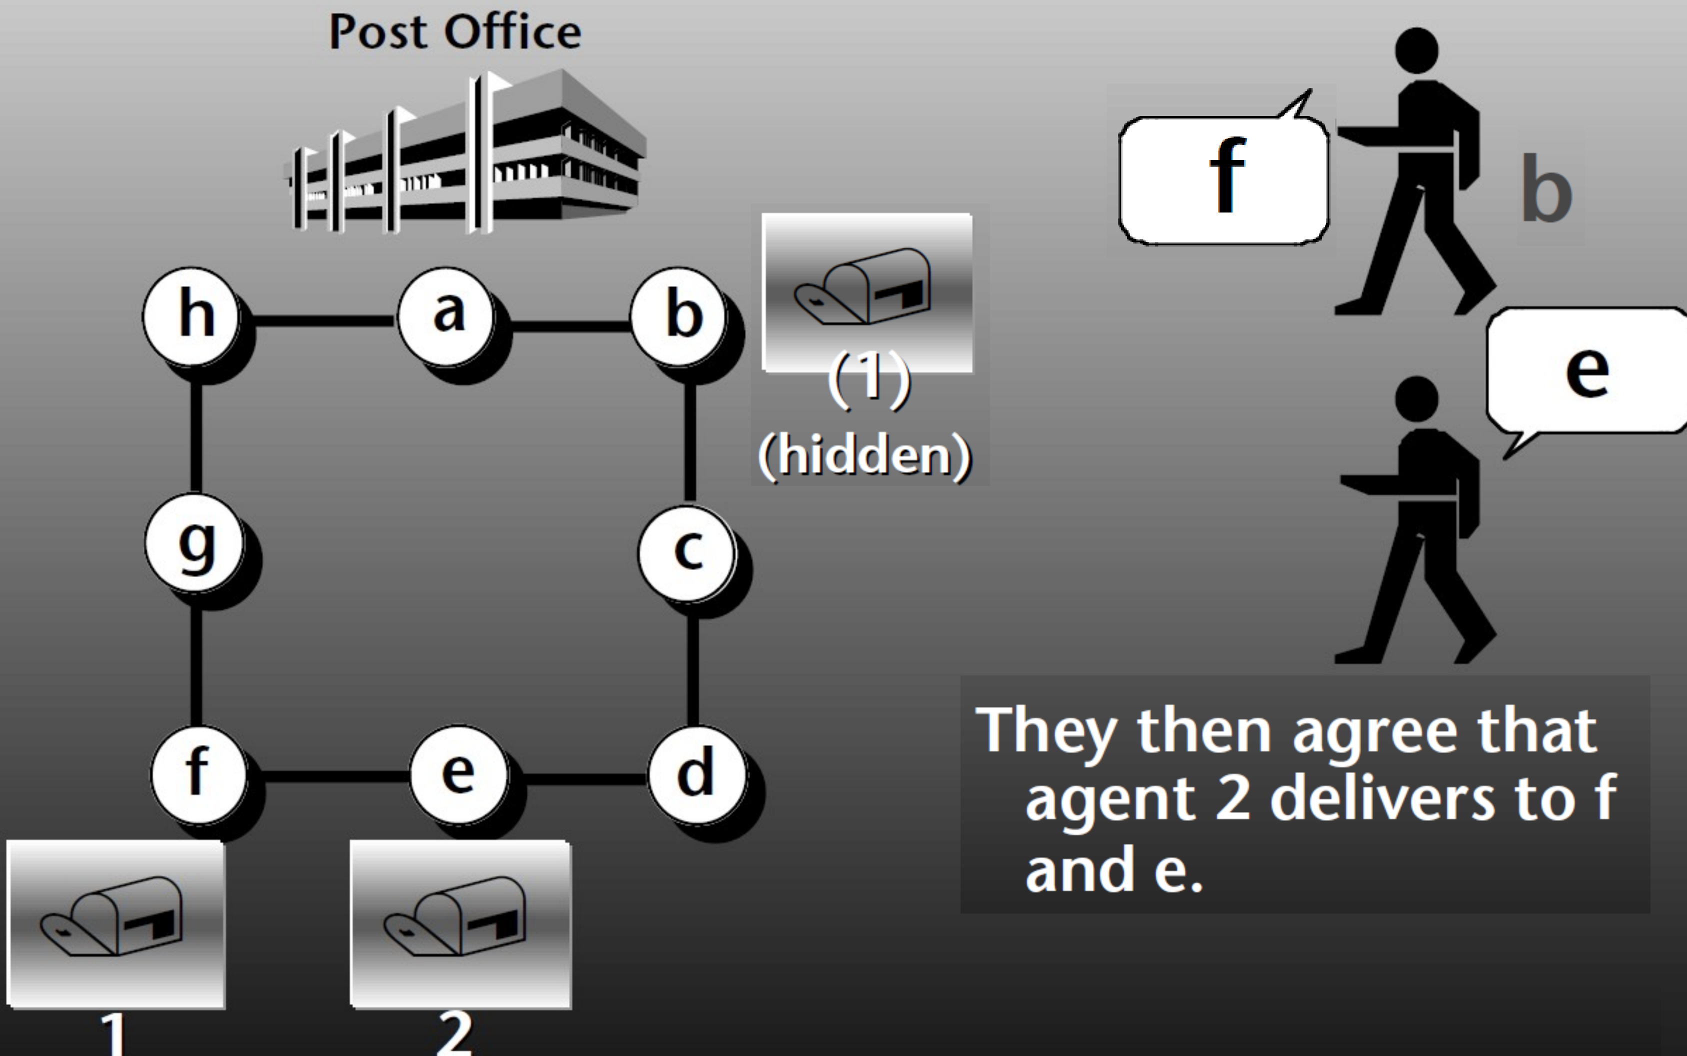

They then agree that agent 2 delivers to f and e.

An Economics View
(Varian 1995)

- "… hyper-rationality may actually be [an] appropriate model for software agents… The whole framework of game theory and mechanism design may well find its most exciting and practical application with computerized agents rather than human agents."

Early Sponsored Search
(Goto 1998)

- Bids are *per click* on a search keyword
- Rank by bid. First price.

pay $12

Autobidders:
Bid minimal amount to maintain current position

Churn…

(Edelman and Ostrovsky, 2007)

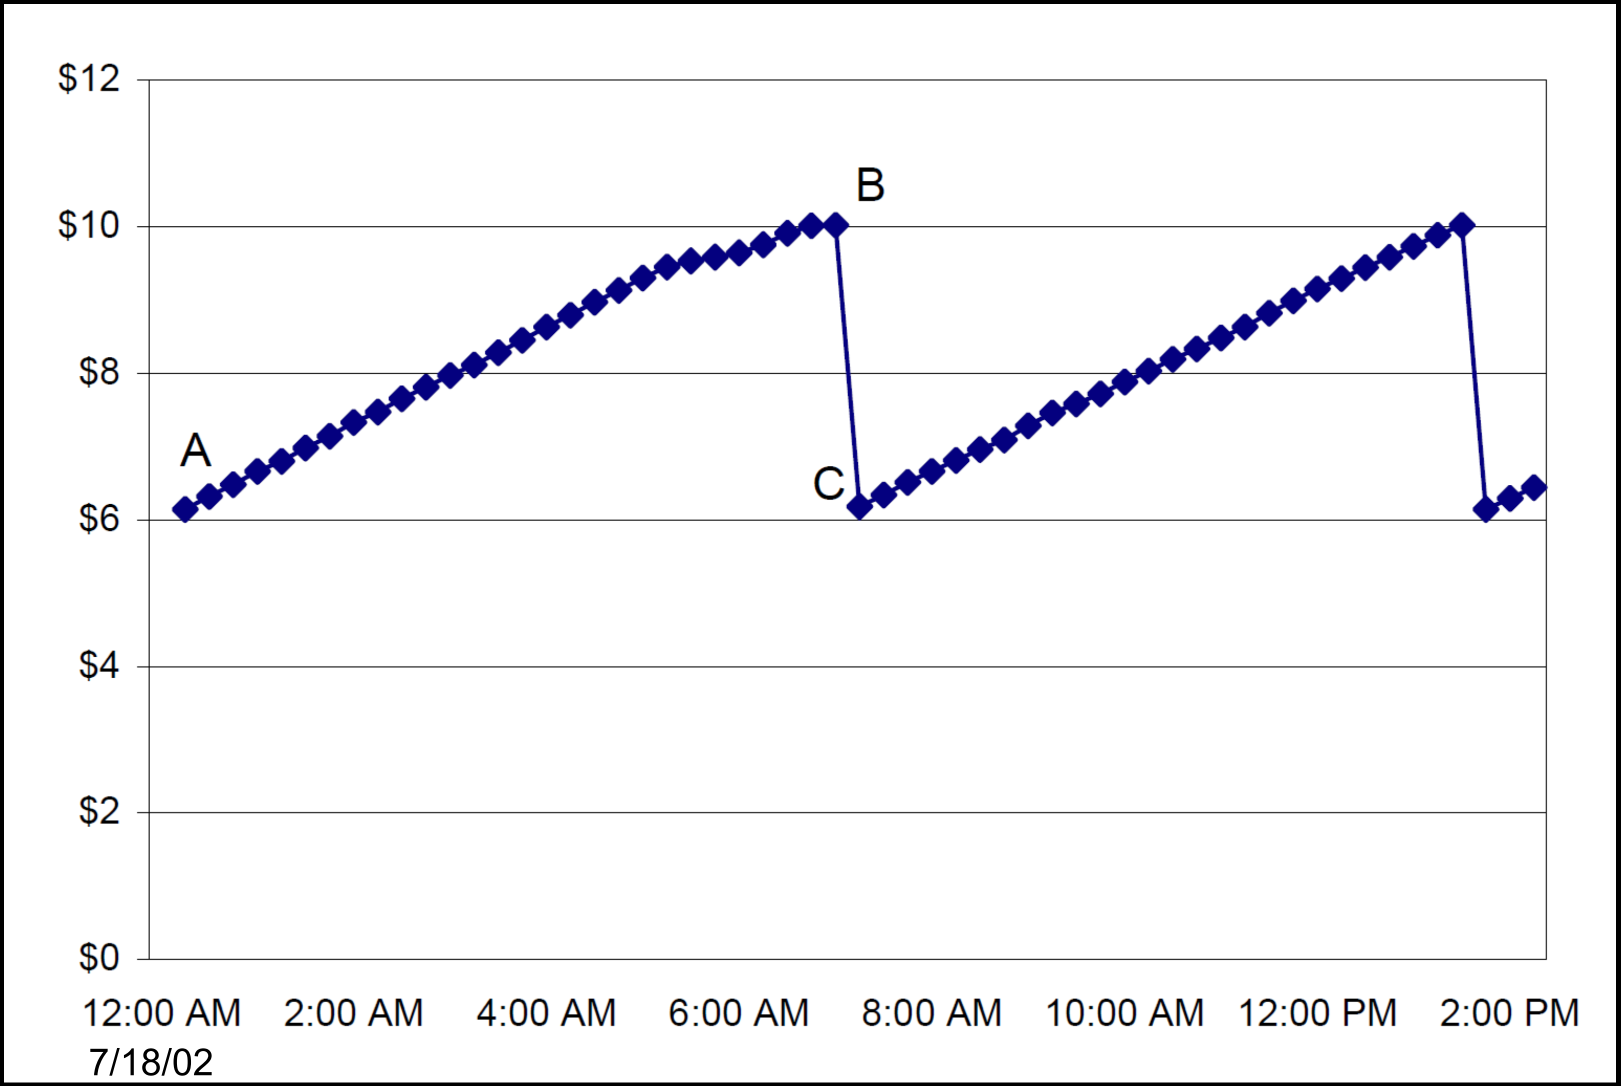

Fix: Generalized Second Price

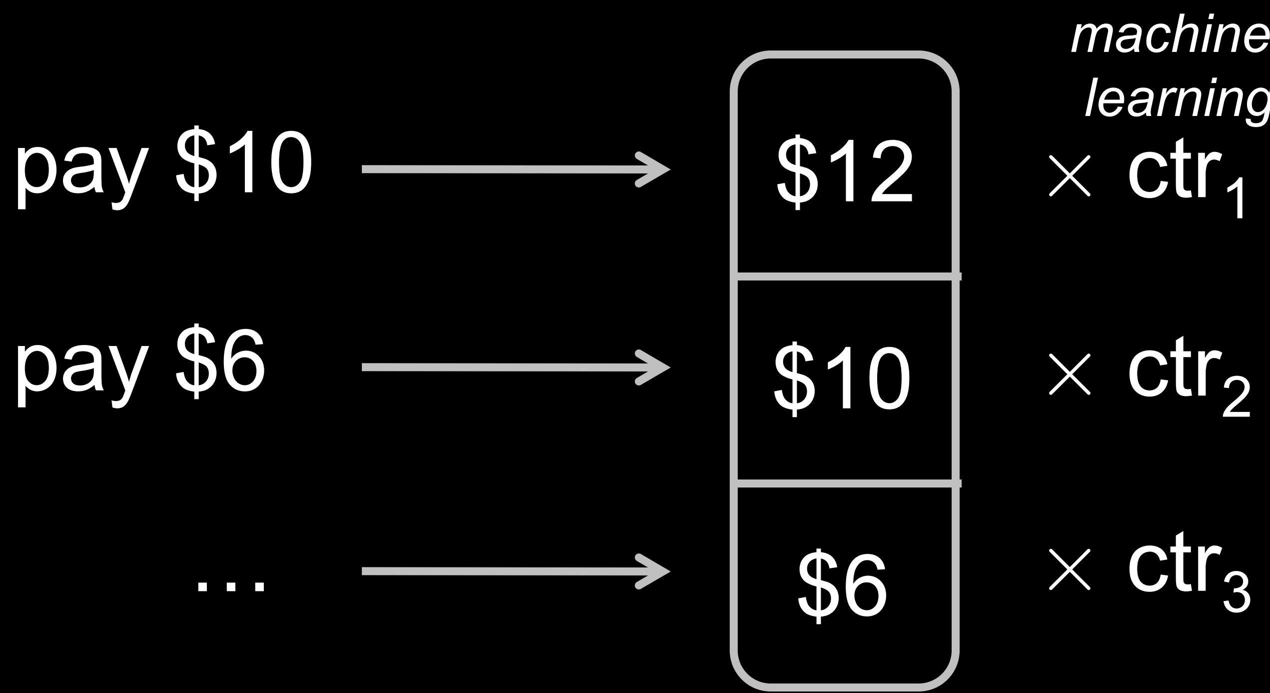

Stability (not full SP)

(1) user relevance, (2) revenue, (3) ad quality

World Design for Self-interested Agents

Mechanism = Algorithm

Example: Combinatorial Auction

(Rassenti, Smith and Bulfin, 1982)

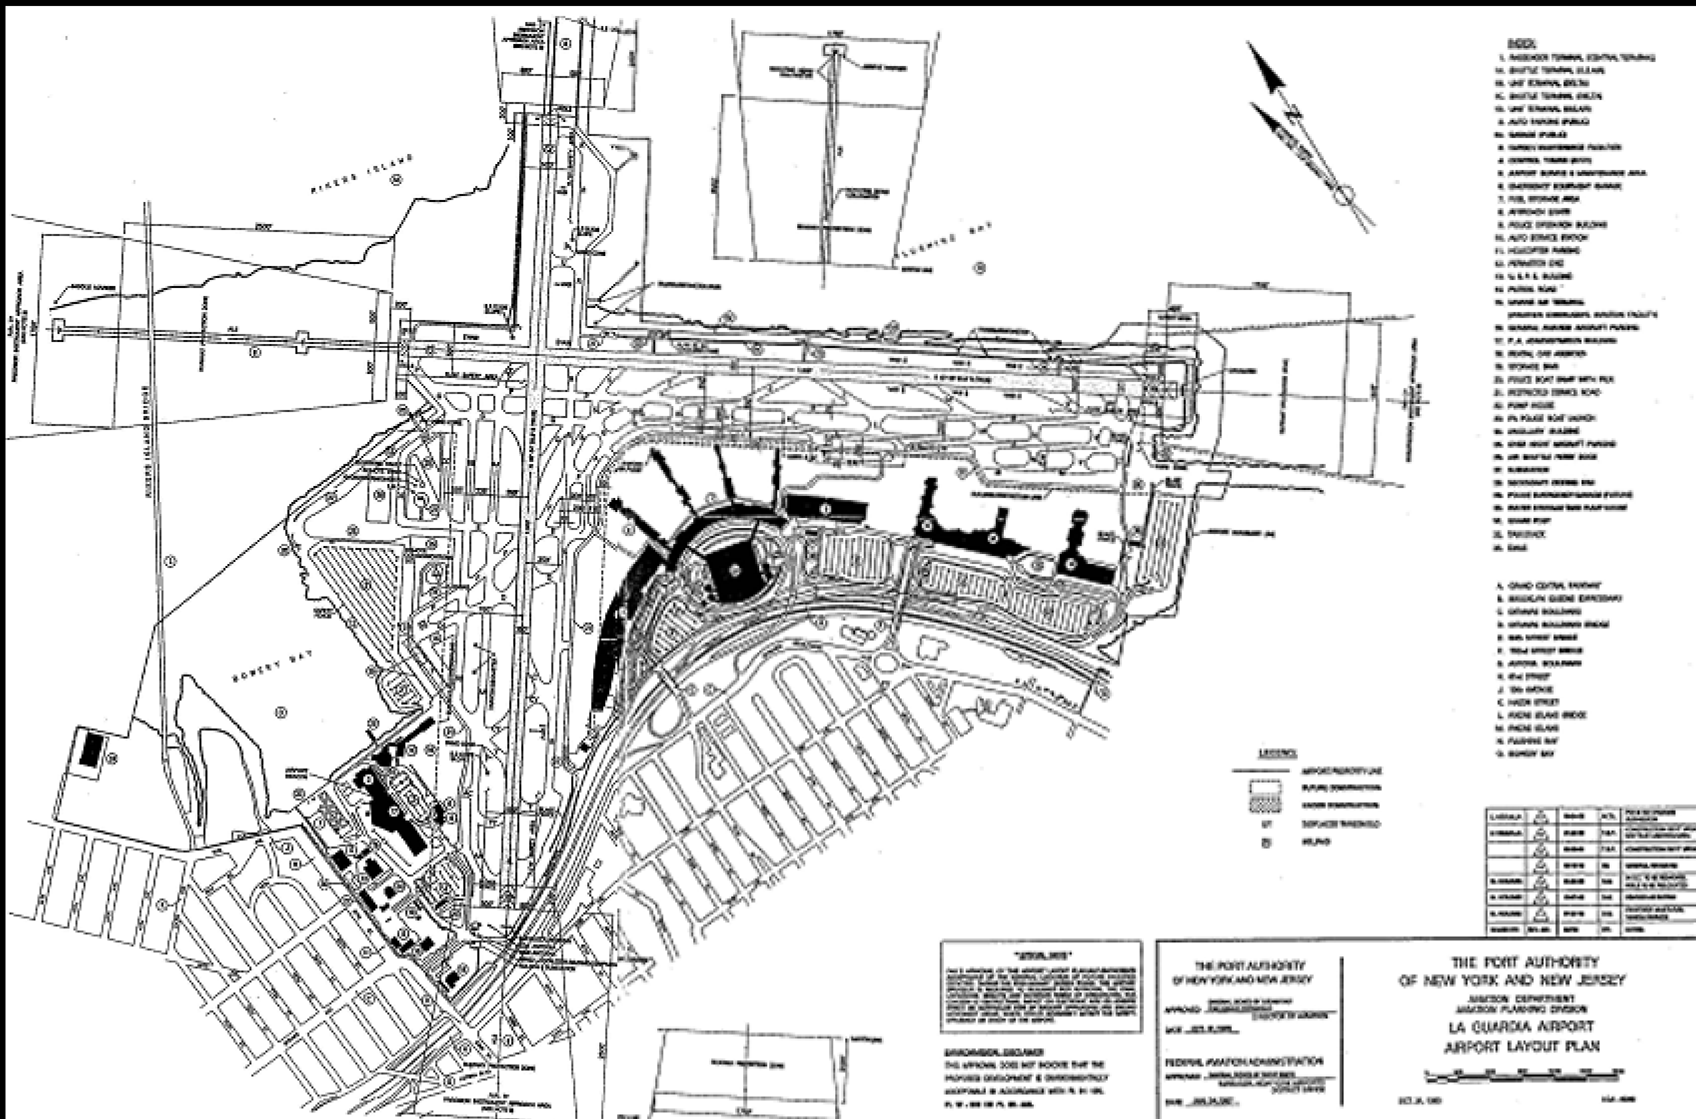

Good Progress

- *Compact and expressive bidding languages*
  - e.g., OR-of-XOR (Sandholm'99), OR* (Fujishima et al.'99, Nisan'00), L_{GB} (Boutilier & Hoos '01)
- *Scalable winner determination*
  - exact algorithms via heuristic search (Fujishima et al.'99, Sandholm'99)
  - tractable special cases (Rothkopf et al.'98)
- *Preference elicitation*
  - iterative CAs (Parkes & Ungar'00), learning theory (Lahaie & Parkes'04), querying (Hudson & Sandholm'03)
  - regret-based methods (Hyafil & Boutilier'06)

An "EconCS" agenda

(Nisan and Ronen'99, Lehmann et al.'99)

- Can't just substitute heuristic algorithms into mechanisms and retain strategyproofness
- Led to a cottage industry in "algorithmic mechanism design"
  - Econ: incentive constraints
  - CS: computational constraints
- Exciting progress

Reasoning about SP mechanisms is hard ☹

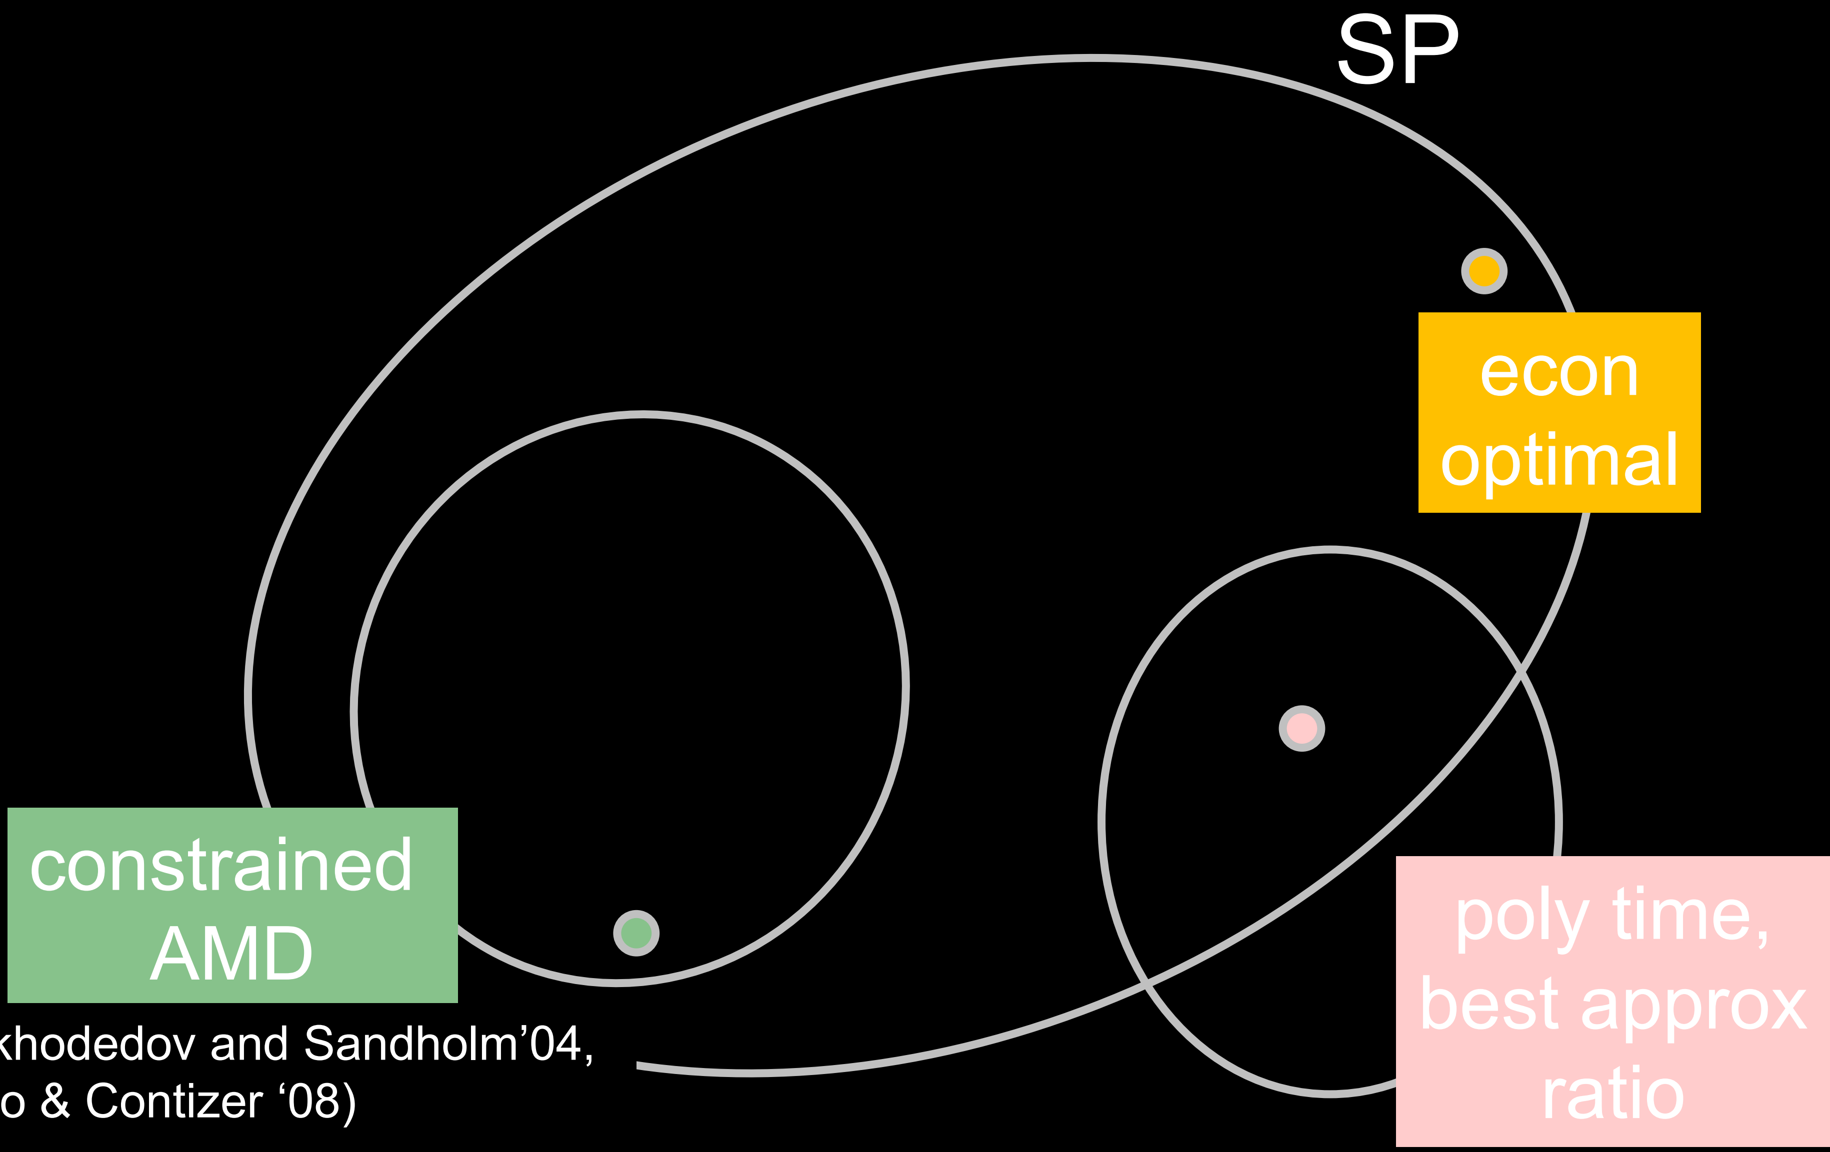

(Likhodedov and Sandholm’04, Guo & Contizer ‘08)

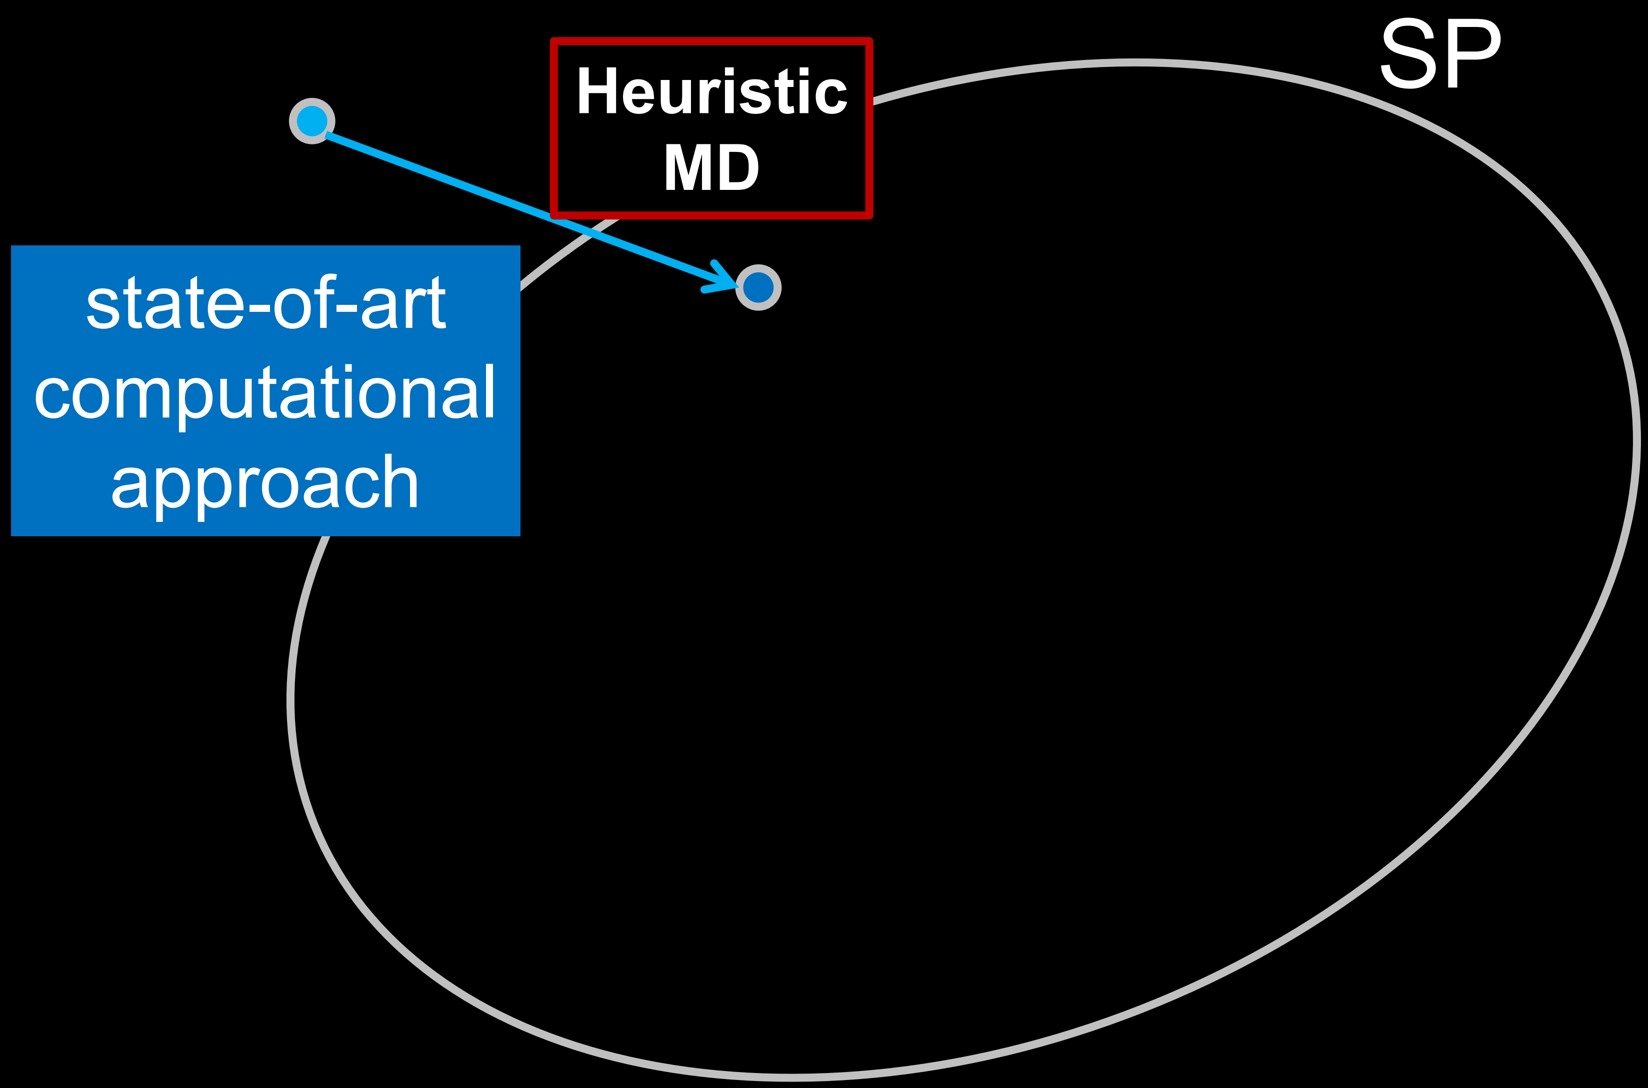
SP
Heuristic
MD
state-of-art
computational
approach

Example: Dynamic Knapsack

m concert tickets to sell. probabilistic model
Agent type: quantity, value, [a,d] interval

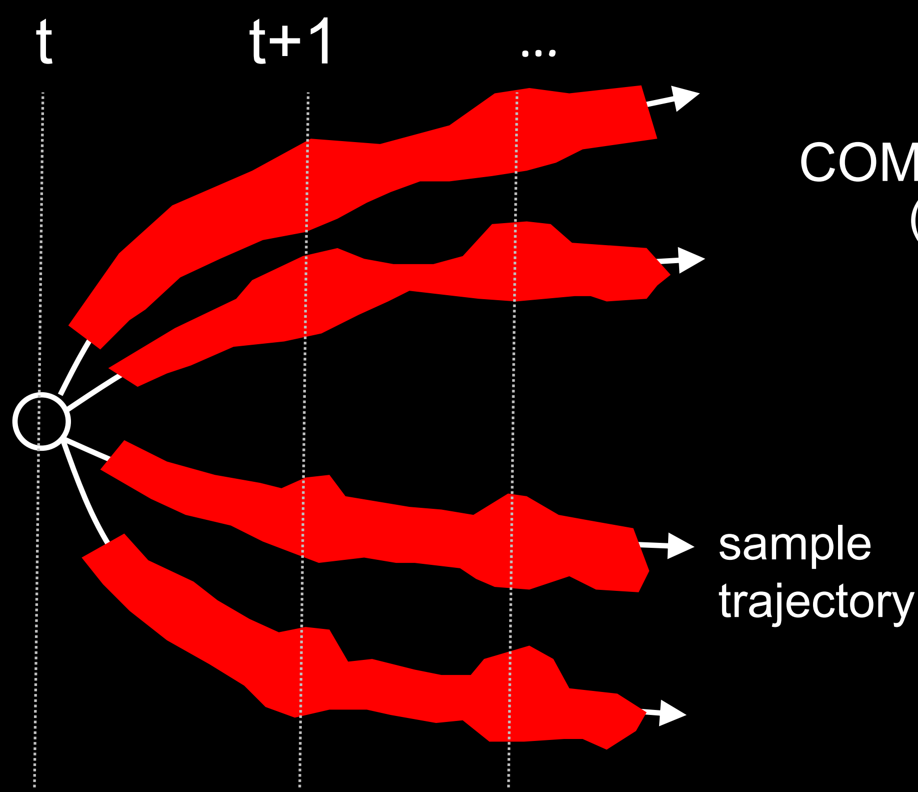

ONLINE STOCHASTIC
COMBINATORIAL OPTIMIZATION
(van Hentenryck and Bent'06)

exogoneous
uncertainty

(P. & Duong '07, Constantin & P.'09)

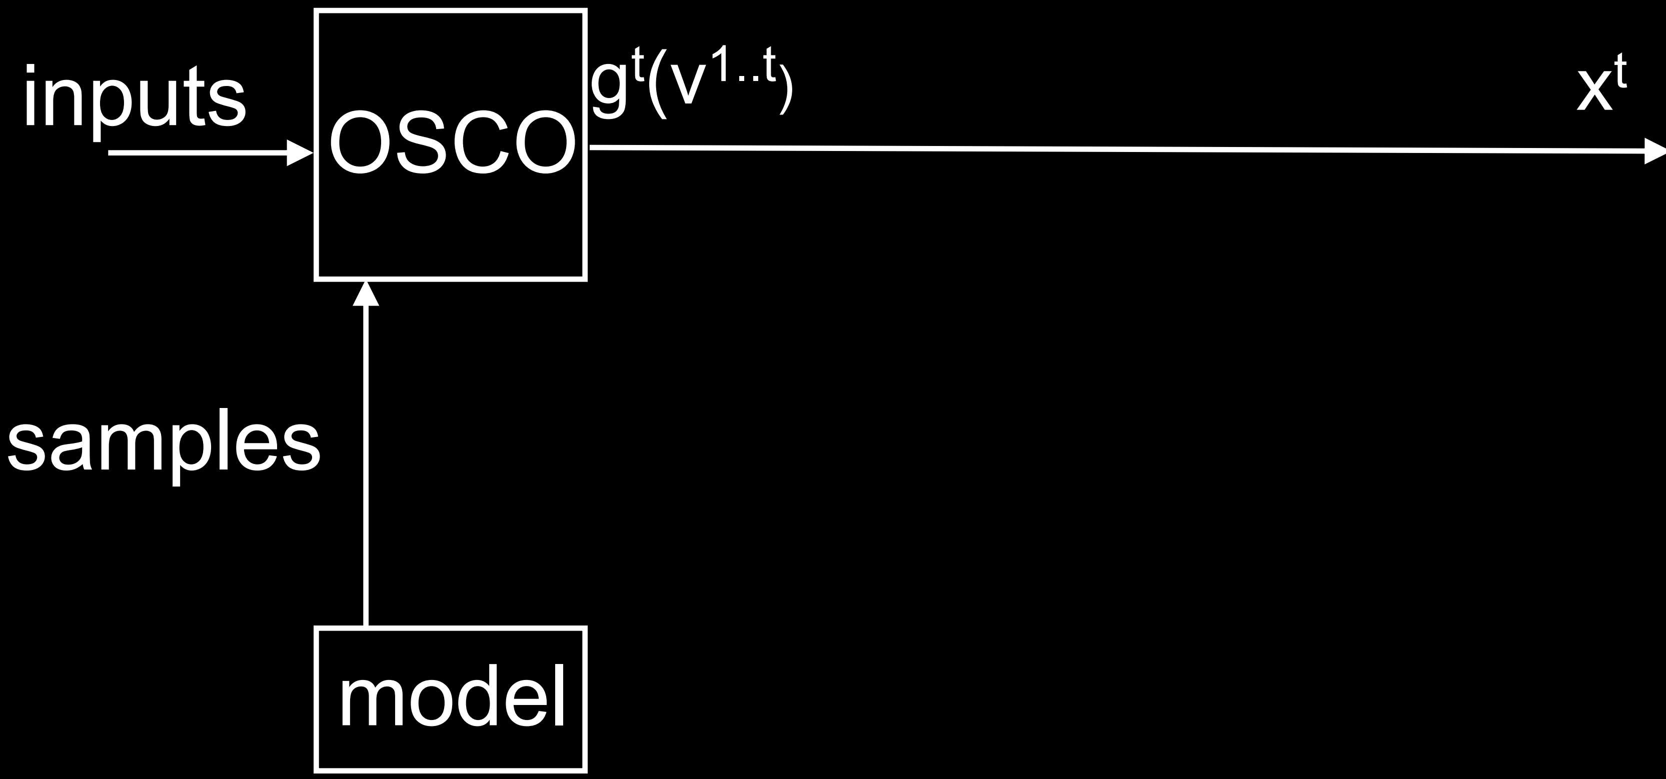

monotonicity

(P. & Duong '07, Constantin & P.'09)

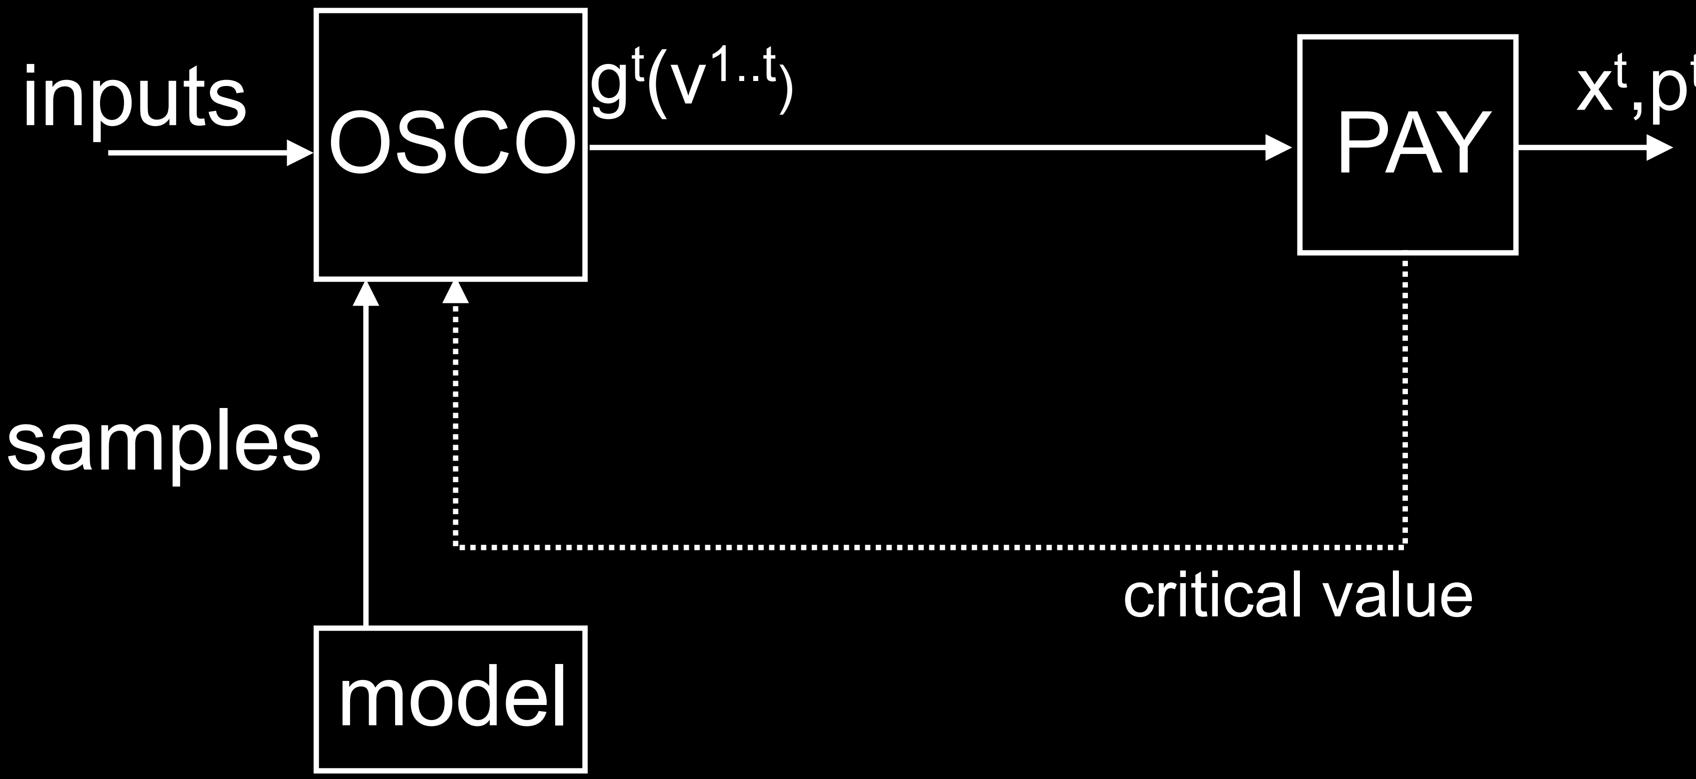

monotonicity

(P. & Duong '07, Constantin & P.'09)

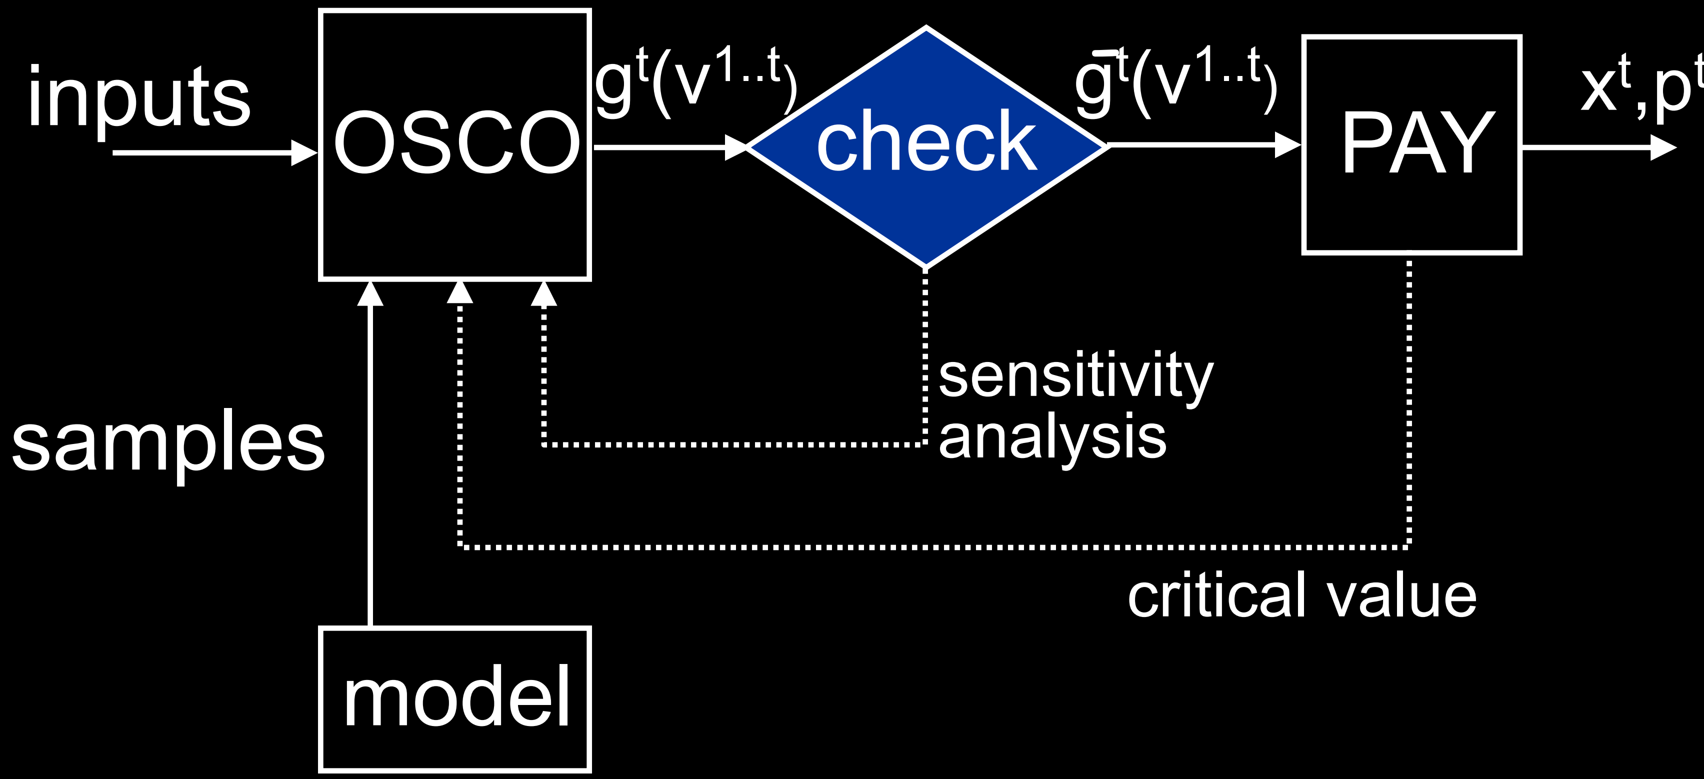

monotonicity

performance (eff):
81.5% best fixed price
89.5% OSCO + ironing
95.2% OSCO

Relaxing away from SP…

- We like SP for reasons of
  - equity (Roth’03, Pathak and Sonmez’08)
  - *simplify reasoning*
  - can predict properties of the mechanism

Relaxing away from SP…

- We like SP for reasons of
  - equity (Roth’03, Pathak and Sonmez’08)
  - *simplify reasoning*
  - can predict properties of the mechanism
- But it is generally hard to obtain
- And, can be provably bad along other dimensions ☹
    - e.g., CAs with complements
      (Ausubel & Milgrom’06, Rastegeri, Condon, & Leyton-Brown’10)

Approx Incentive Alignment

- A satisfactory answer will:
  - allow for comparison of mechanisms
  - allow for a larger design space
  - still provide predictable behavior

Old Favorite: Min Max Regret

- Regret = best utility – actual utility
- Maximally SP: minimizes max regret across agents on every instance
- ϵ-SP: max regret $\leq \epsilon$

Example: Comb. Exchange

- Airlines buying and selling landing slots

Example: Comb. Exchange

- Airlines buying and selling landing slots

- $p_{vcg,i}$ = bid – marginal contribution
  (Δ_{vcg})

Example: Comb. Exchange

- Airlines buying and selling landing slots

- $p_{vcg,i}$ = bid – marginal contribution
  (Δ_{vcg})

- Runs at a deficit in a CE ☹
- Impose $\sum p_i \geq 0$

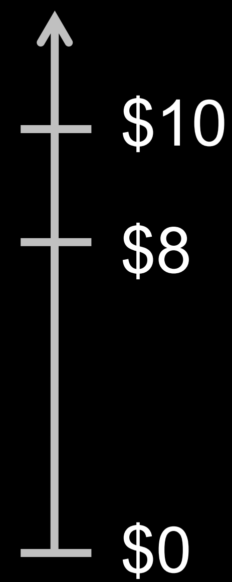

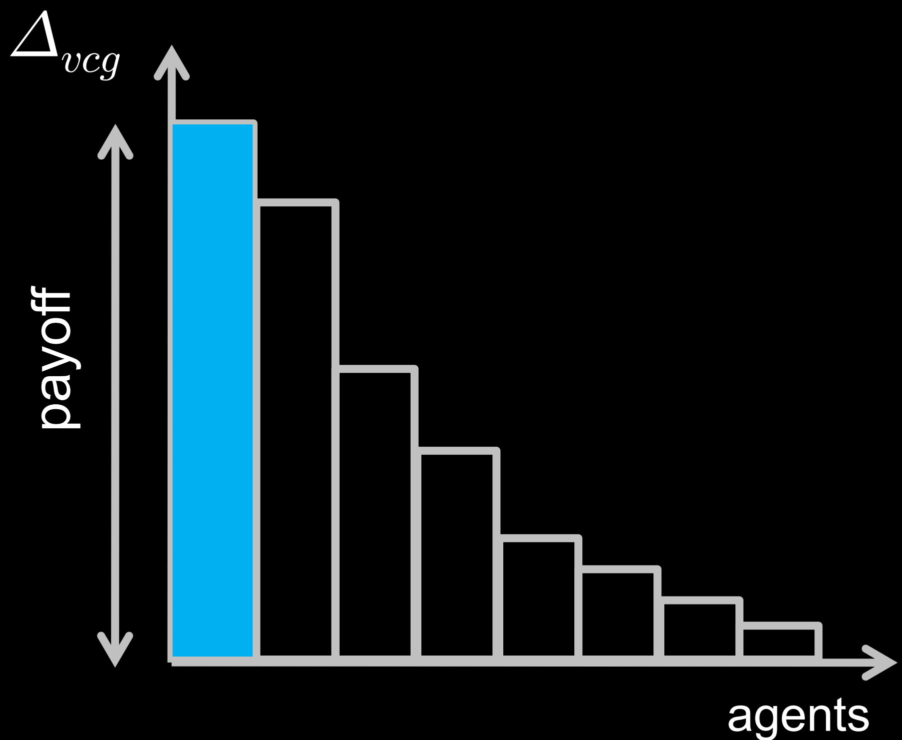
Δ_{vcg}
payoff
agents

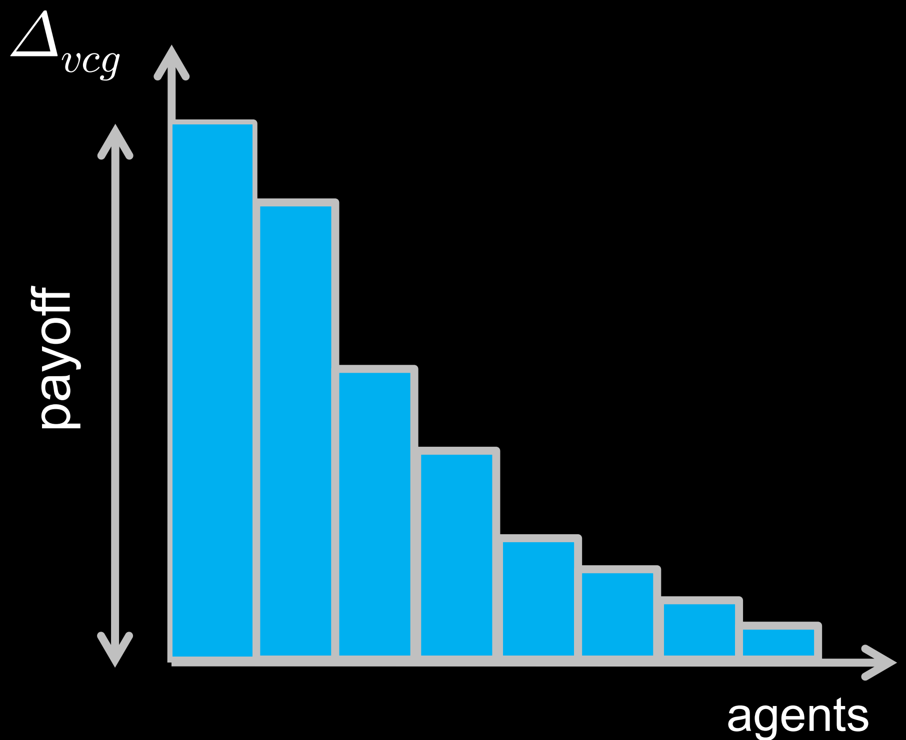

Δ_{vcg}
payoff
agents

Two mechanism rules

(Parkes, Kalagnanam and Eso '01)

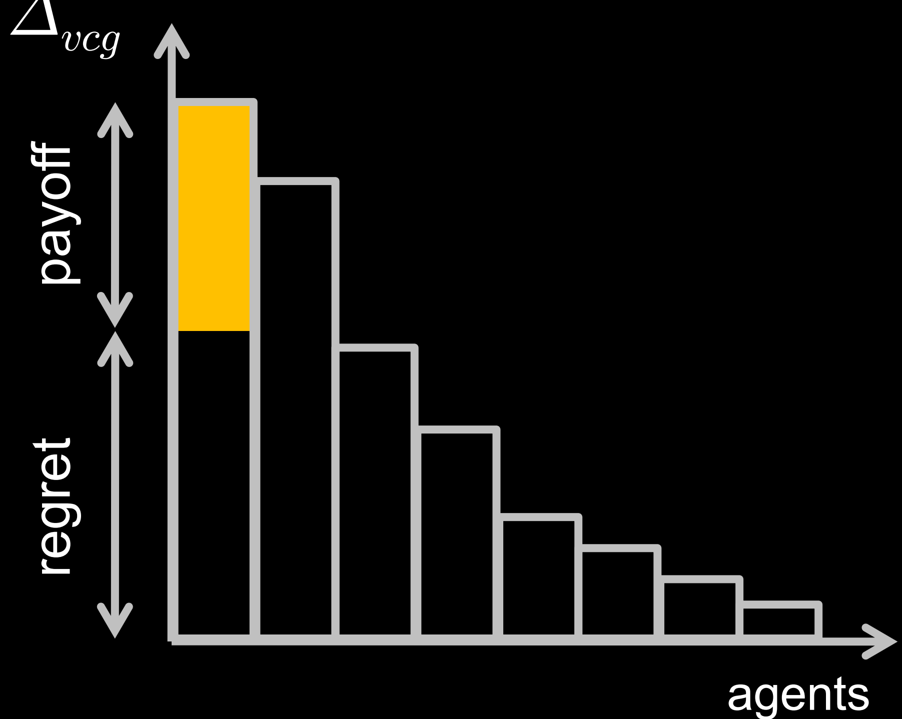

Two mechanism rules

(Parkes, Kalagnanam and Eso '01)

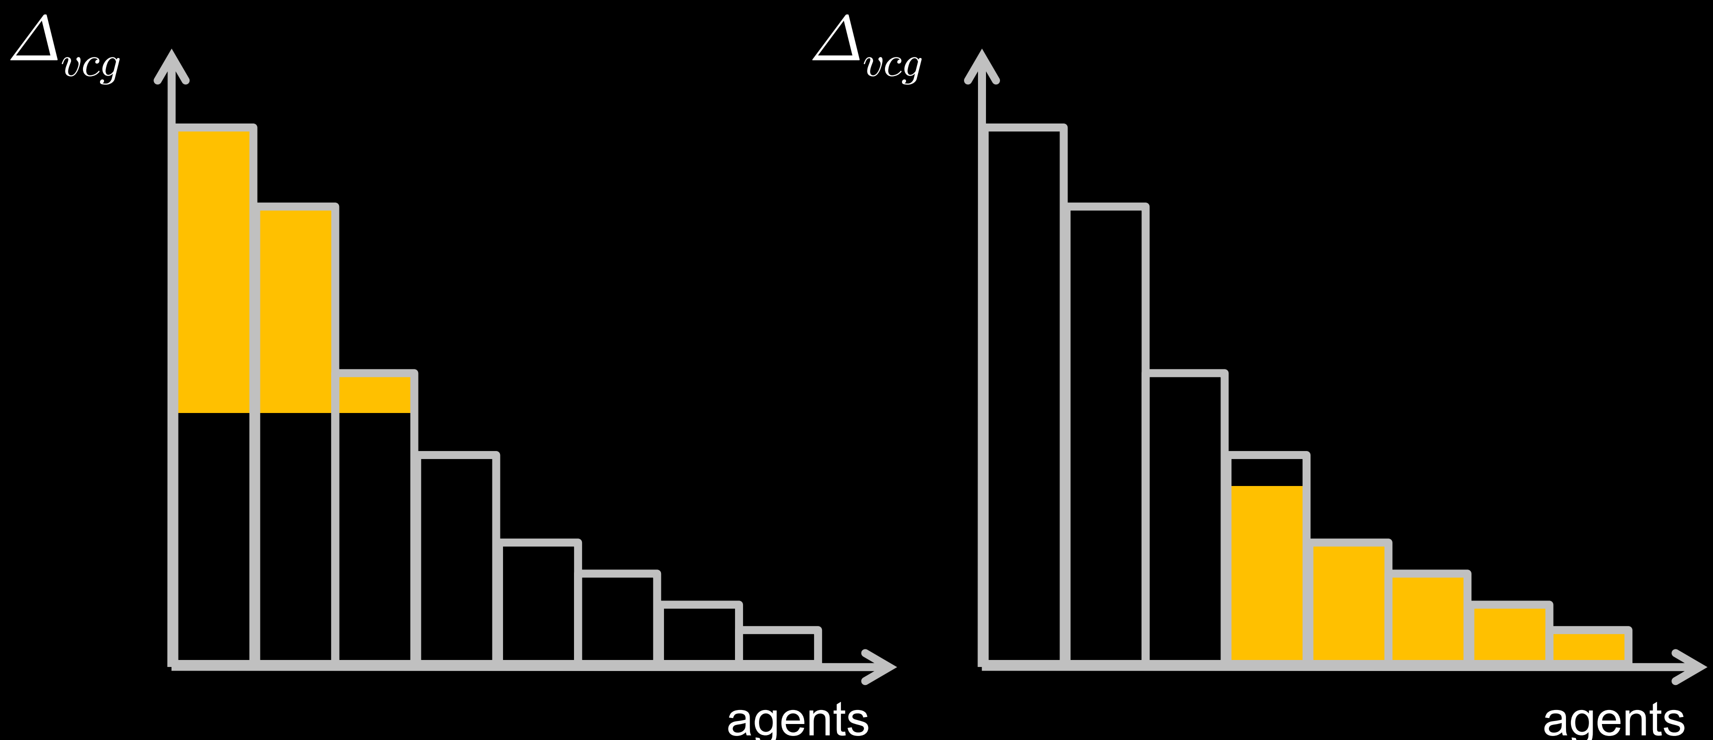

Threshold rule
(min max regret)

Small rule
max #(regret=0)

Approximate BNE Analysis

(Lubin & Parkes ’09)

	strategy			efficiency		
Rule	**Dec.**	**Uni.**	**Sup.**	**Dec.**	**Uni.**	**Sup.**
VCG	*0.0*	*0.0*	*0.0*	*100*	*100*	*100*
Two Triangle	**0.1**	**0.4**	**5.6**	**99.99**	**100**	**97.95**
Threshold	14.6	27.2	11.2	93.64	81.09	89.74
Small	**0.0**	**0.1**	**0.2**	**99.99**	**100**	**100**
No Discount	62.3	80.9	72.4	34.15	50.11	48.21

(For BNE, see Vorobeychik & Wellman’08, Rabinovich, Gerding, Polukarov & Jennings’09)

Distributional View: Payoffs

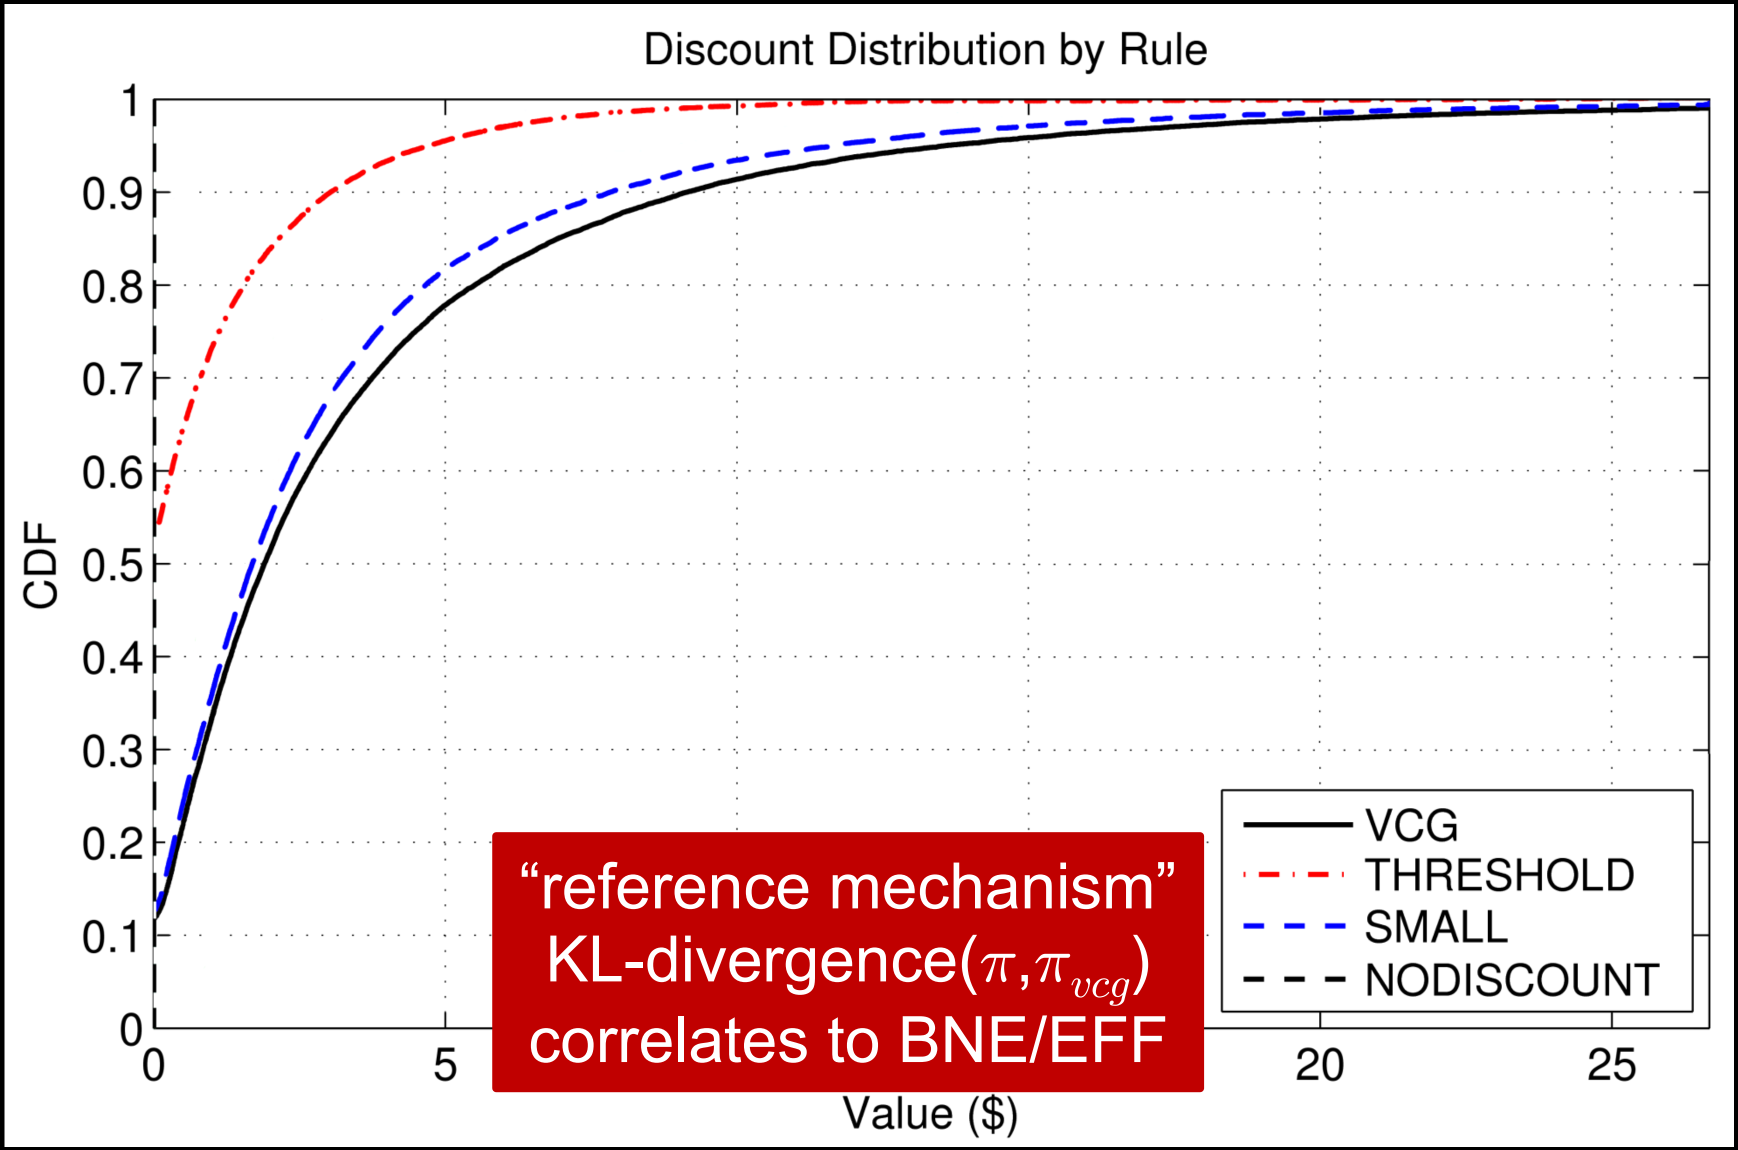

Regret Quantiles

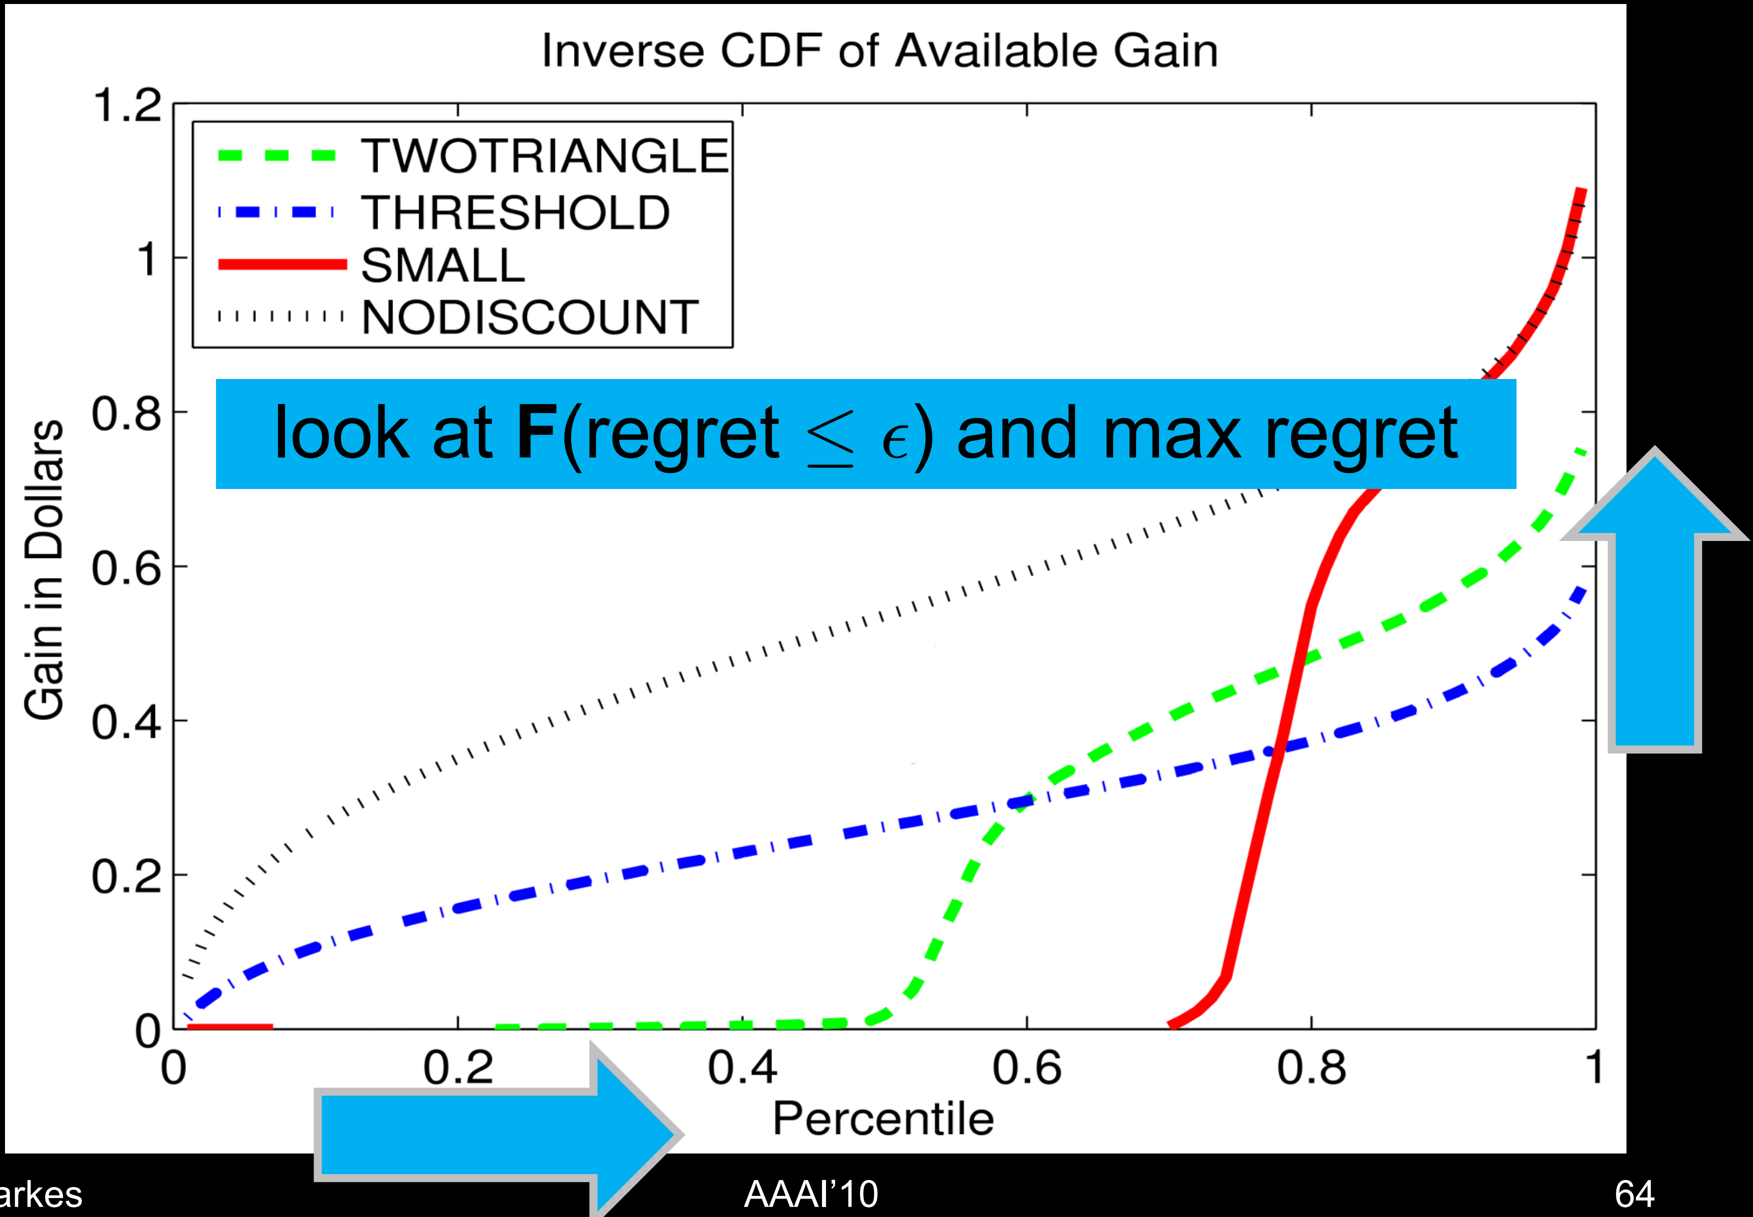

look at **F**(regret $\leq \epsilon$) and max regret

From Events to Platforms

- eBay, sponsored search, display advertising on Facebook, etc. are all *dynamic* problems:
  - Dynamic population
  - Learning by agents
  - Learning by the mechanism
  - Uncertain supply
- Need incentive engineering to coordinate "always on" dynamic systems

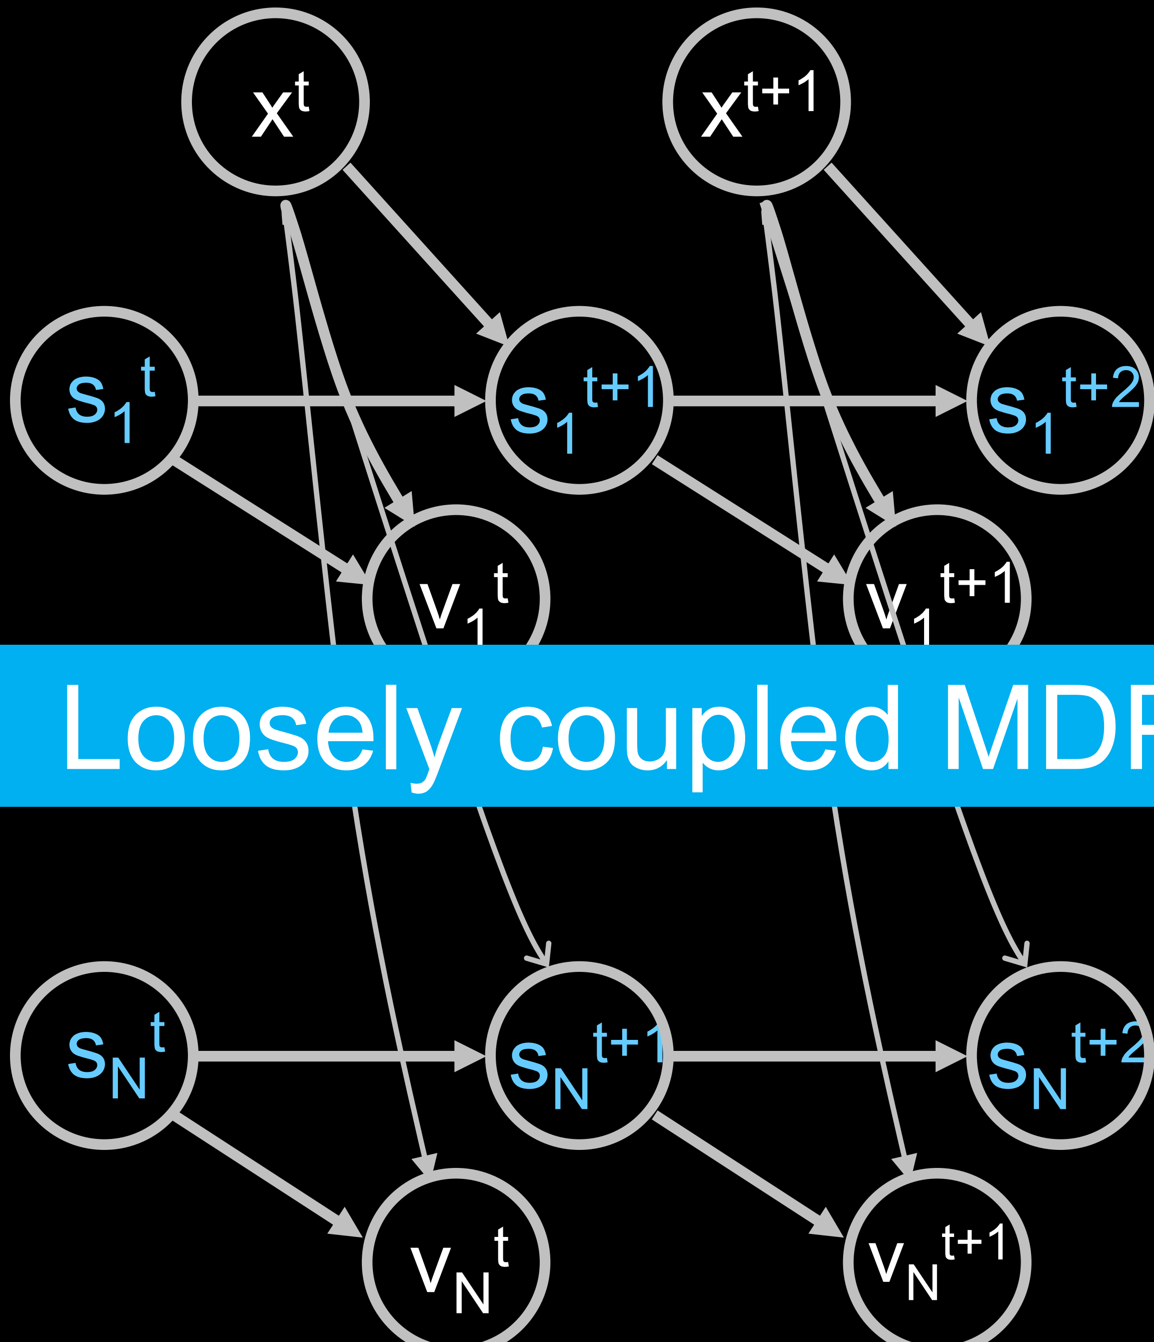

Loosely coupled MDPs

Theory: Dynamic VCG

- Support Optimal MDP policies
- With dynamic population, static types
  - *includes dynamic Cas*
  - P. & Singh '03, P., Singh & Yanovsky'04
- With static population, dynamic types
  - *includes Bayesian optimal learning*
  - Bergemann & Valimaki '08
- Unified view
  - Cavallo, P. & Singh'09

Skill Acquisition Example

(Cavallo & Parkes'08)

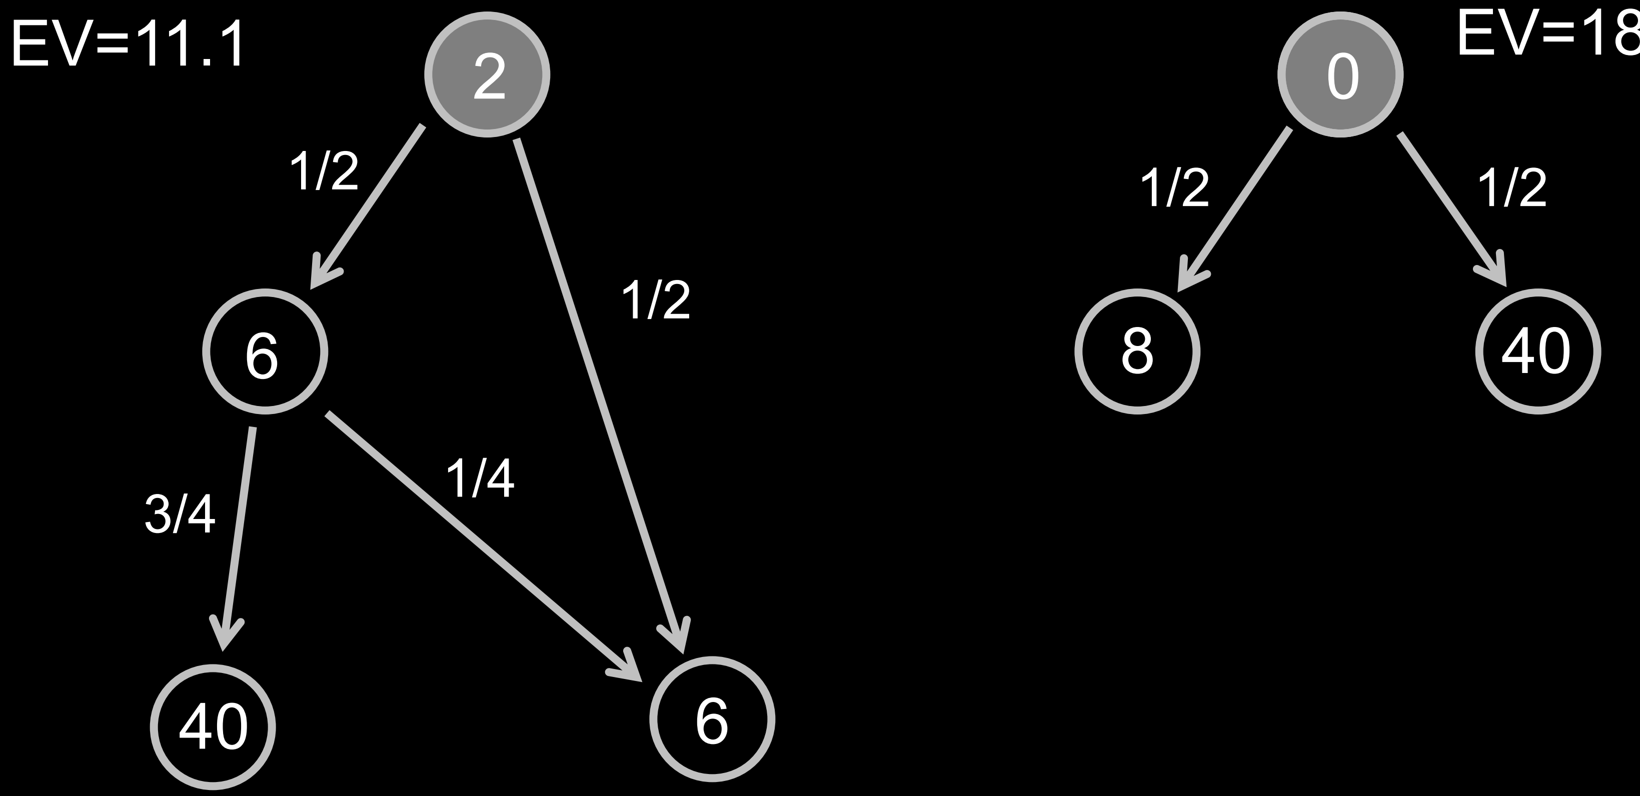

$\beta = 0.75$

Skill Acquisition Example

(Cavallo & Parkes'08)

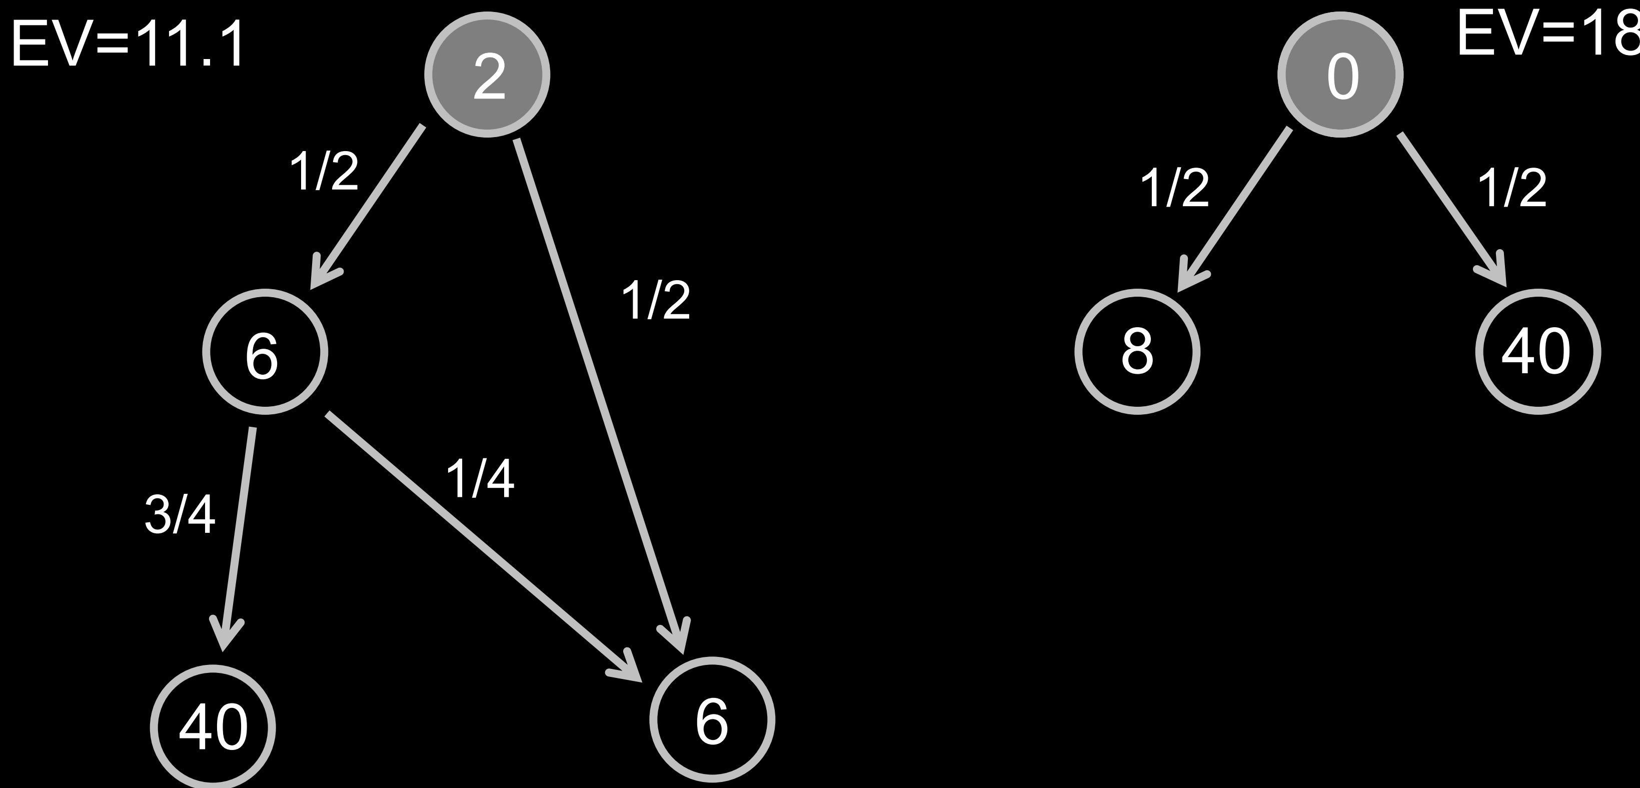

$\beta = 0.75$

Skill Acquisition Example

(Cavallo & Parkes'08)

payment
$(1-\beta)$ 11.1 = 2.78

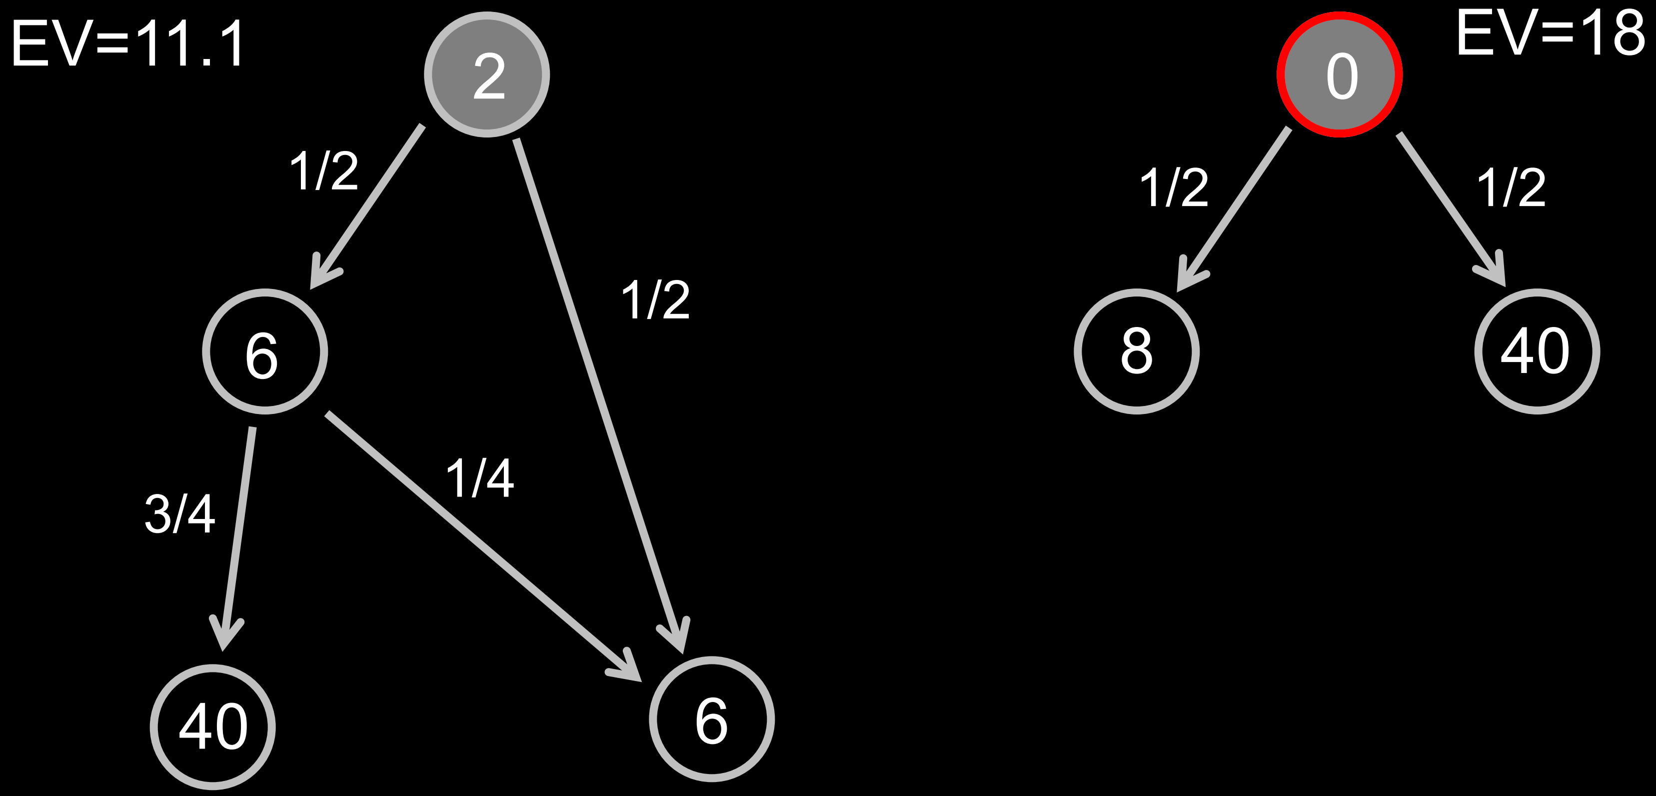

$\beta = 0.75$

Skill Acquisition Example

(Cavallo & Parkes'08)

payment

$(1-\beta)$ 11.1 = 2.78

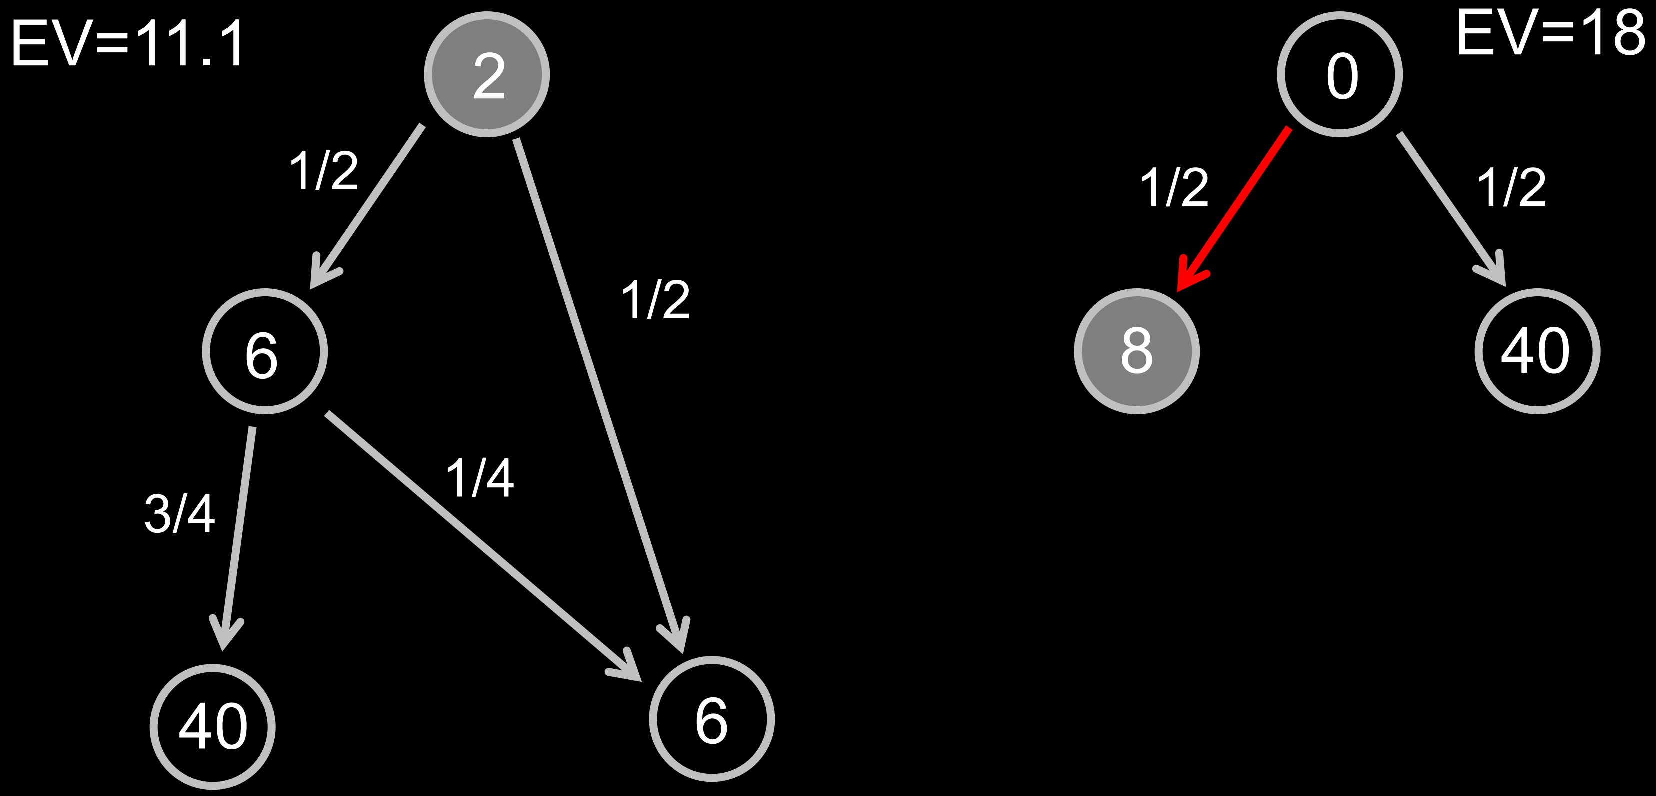

$\beta = 0.75$

Skill Acquisition Example

(Cavallo & Parkes'08)

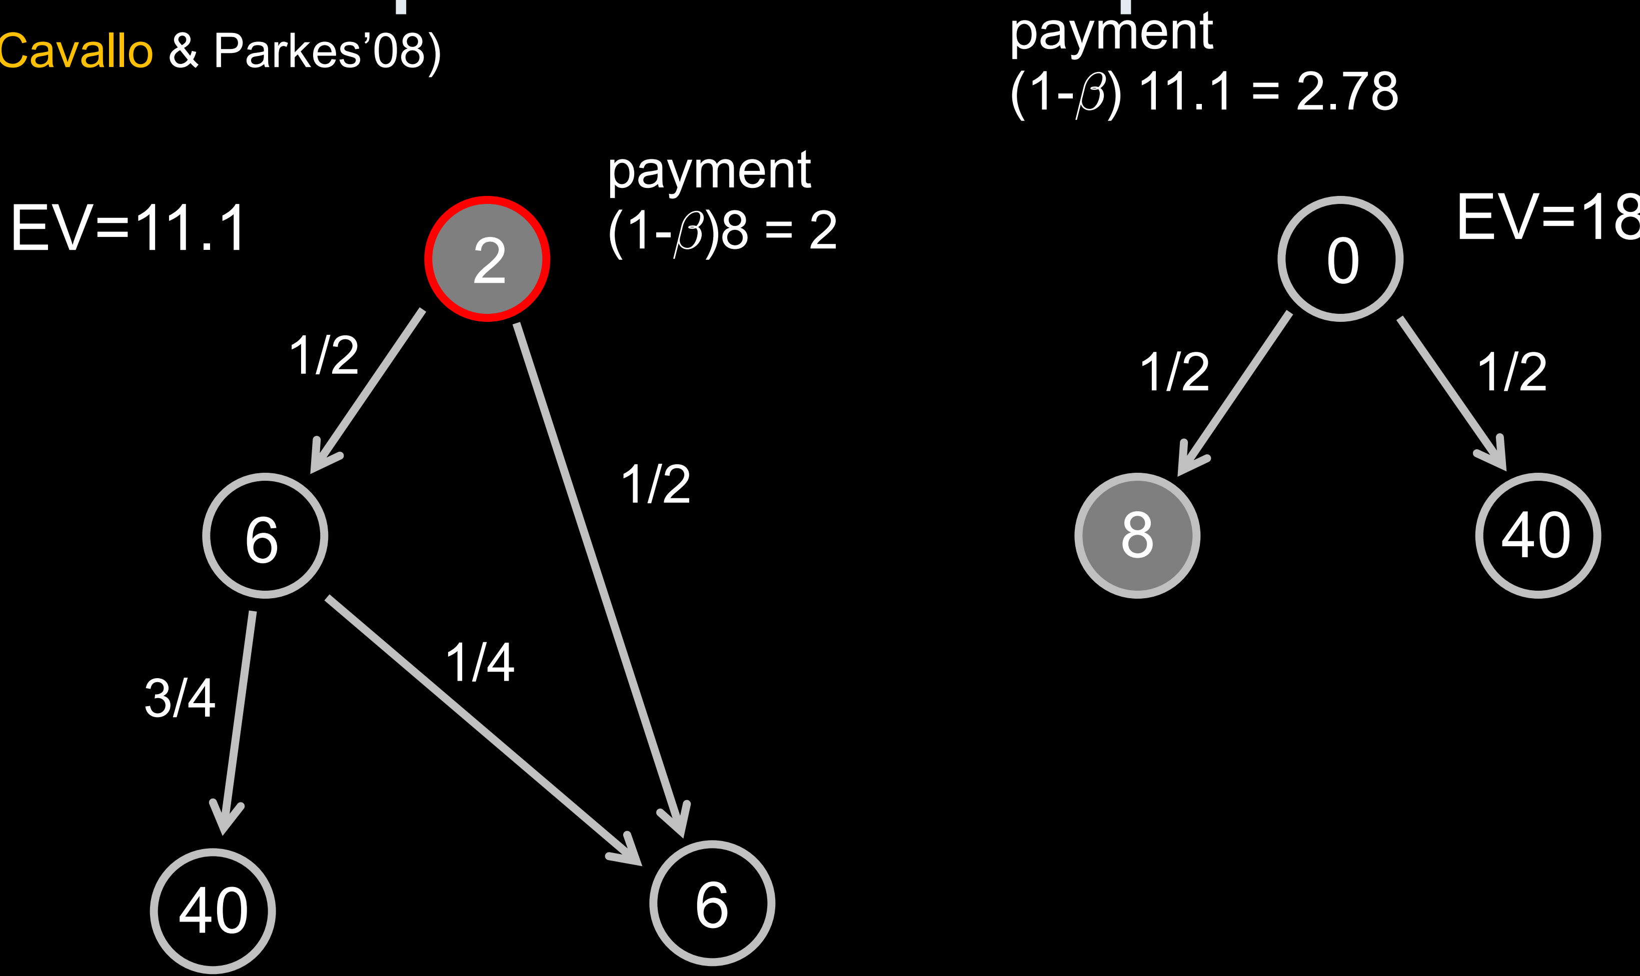

$\beta = 0.75$

Skill Acquisition Example

(Cavallo & Parkes'08)

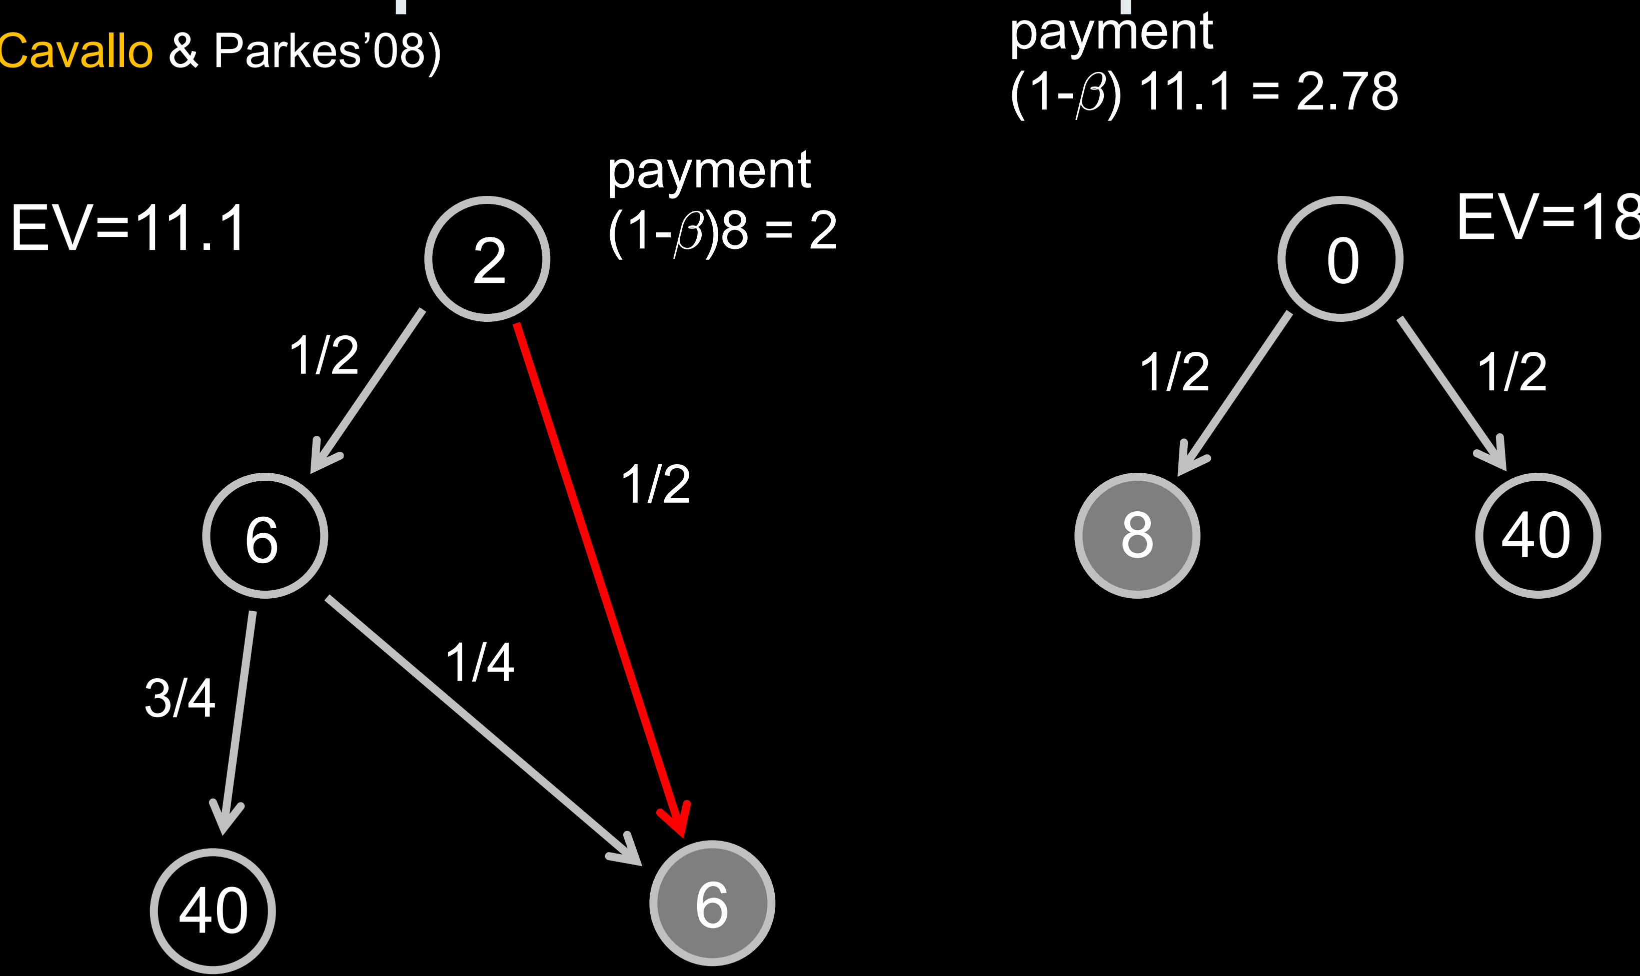

β = 0.75

Skill Acquisition Example

(Cavallo & Parkes'08)

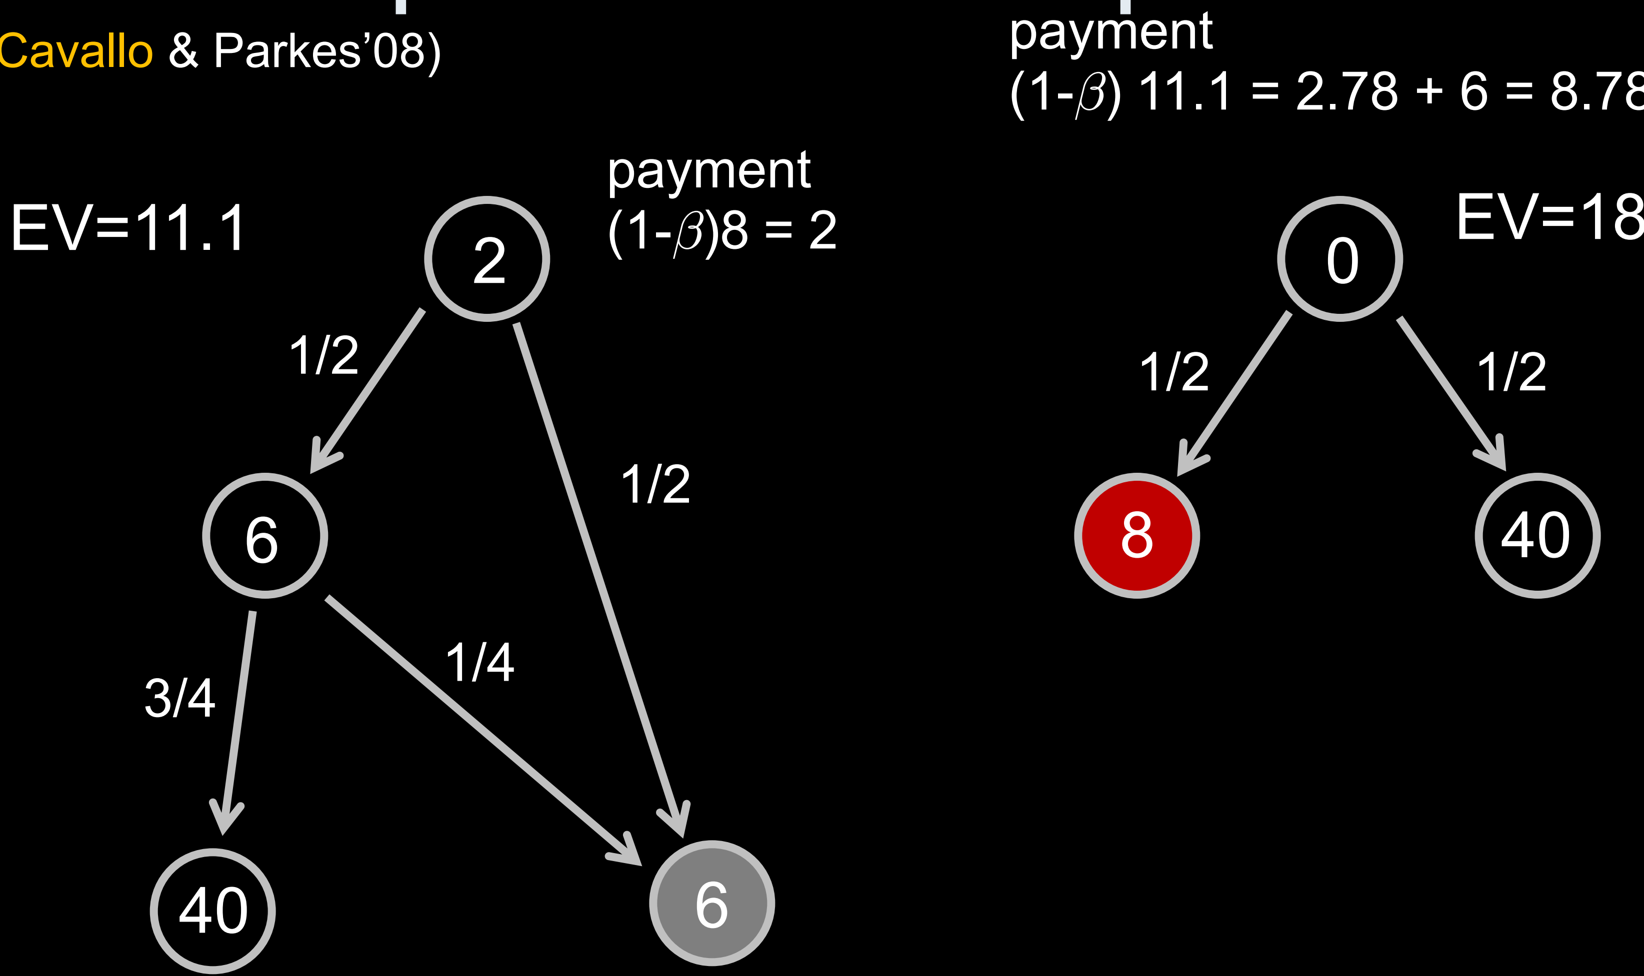

$\beta = 0.75$

Dynamic-VCG: Scaling-up

- Need optimal-in-range policies

$$\pi^* \in \arg\max_{\pi \in \Pi'} V^{\pi}(s)$$

$\Rightarrow$ an interesting *meta*-problem

(see Gerding, Stein, Larson, Rogers & Jennings'10)

Back to tasks…

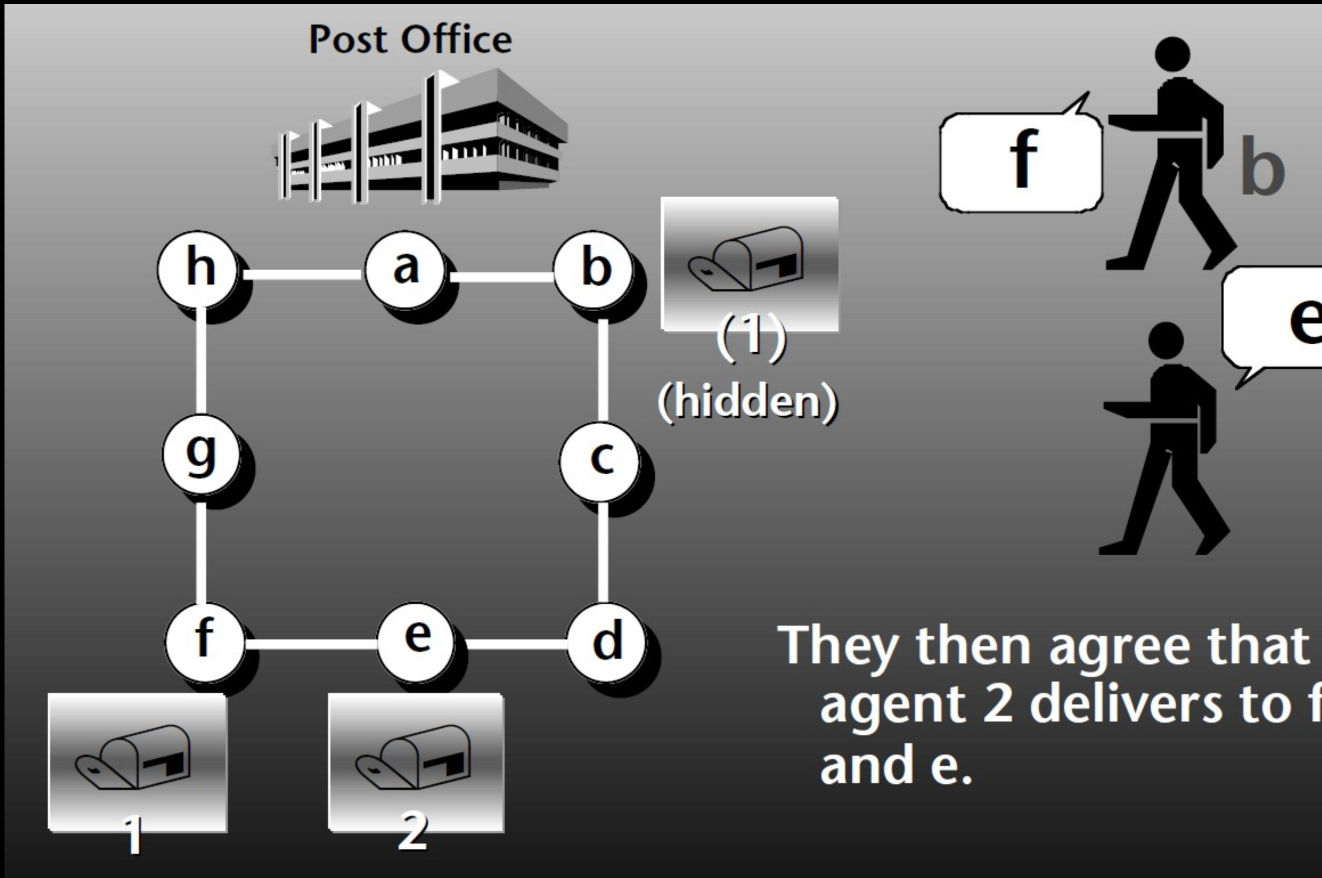

They then agree that agent 2 delivers to f and e.

Crowdsourcing Platforms

- *Amazon Mechanical Turk*
  - online labor market for “human intelligence tasks” (e.g., data cleaning)
- *InnoCentive* (innovation marketplace)
  - 150+ seekers, 180,000+ solvers, $$ prizes
  - 900+ challenges
  - e.g., “Sustainable Packaging for Developing World”
- *TopCoder* (code development)
  - 250,000+ workers, $$ to first and second-best
  - e.g., NASA/HBS/LBS “MedKit optimization”

The Longitude Prize

http://www.nmm.ac.uk/harrison

- Royal Observatory
  - founded in 1675 to solve the "longitude problem"
  - sailors could measure local time from sun, with an accurate reference time, could compute longitude

The Longitude Prize

http://www.nmm.ac.uk/harrison

- Royal Observatory
  - founded in 1675 to solve the "longitude problem"
  - sailors could measure local time from sun, with an accurate reference time, could compute longitude
- Won by John Harrison (1693-1776)
  - started work in 1730, awarded prize at age 79 in 1773

But rapid integration of partial solutions from multiple sources is new

NetFlix Prize

www.the-ensemble.com

$y(x) = g(y_1(x), \ldots, y_K(x))$
meta-features

~volinsky/netflix/bpc.html

Gravity

Dinosaur Planet

(Fall 2007)
When Gravity and Dinosaurs Unite

Grand Prize Team

(Jan 2009, share 2/3 prize for final 1% improvement)

+0.21% Bertino

+0.14% Sill

+0.06% Nabutovsky

+0.08% Sill

+0.19%

Opera Solutions

Vandelay Industries !

June 26, 2009

The Ensemble +0.43%

=10.10% improvement
July 26, 2009

DARPA “Red Balloons”

- Launched Oct 29, 2009.
- Ten 8’ red balloons, 30.5 m in air
- $40,000 prize (for latitude and longitude)
- Competition @ 10am, December 5, 2009

DARPA "Red Balloons"

- Launched Oct 29, 2009.
- Ten 8' red balloons, 30.5 m in air
- $40,000 prize (for latitude and longitude)
- Competition @ 10am, December 5, 2009
- Won by 6:52pm!

MIT: Recursive Incentive Scheme

Recruited 5,400 individuals in 36 hours
One-time “supply chain”

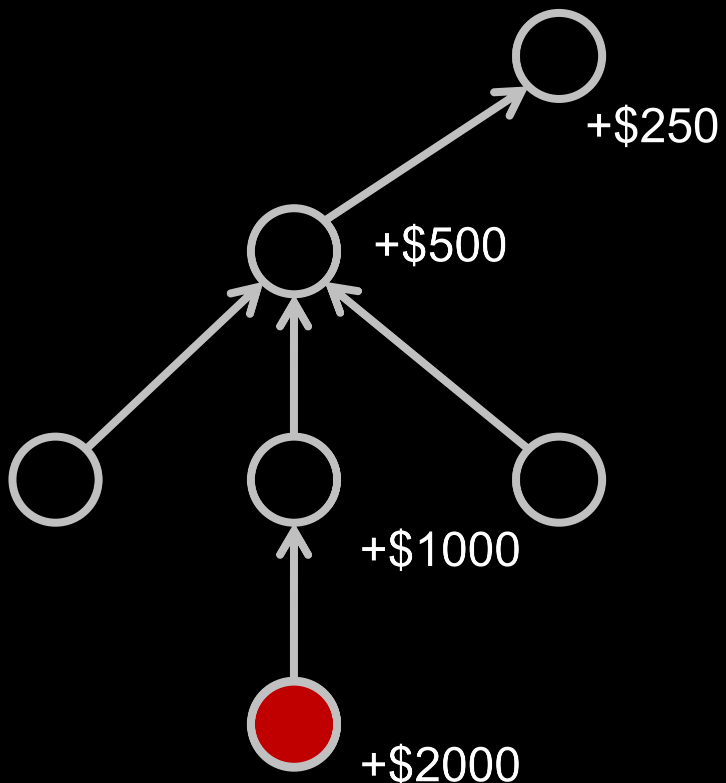

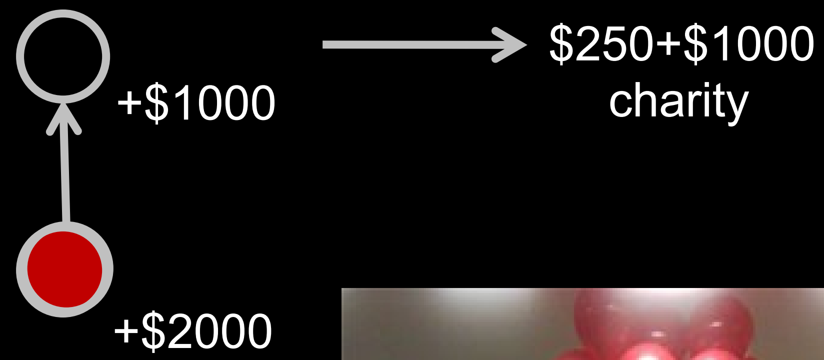

Anmol Madan, Galen Pickard, Riley Crane
Alex ("Sandy") Pentland, Wei Pan
and Manuel Cebrian

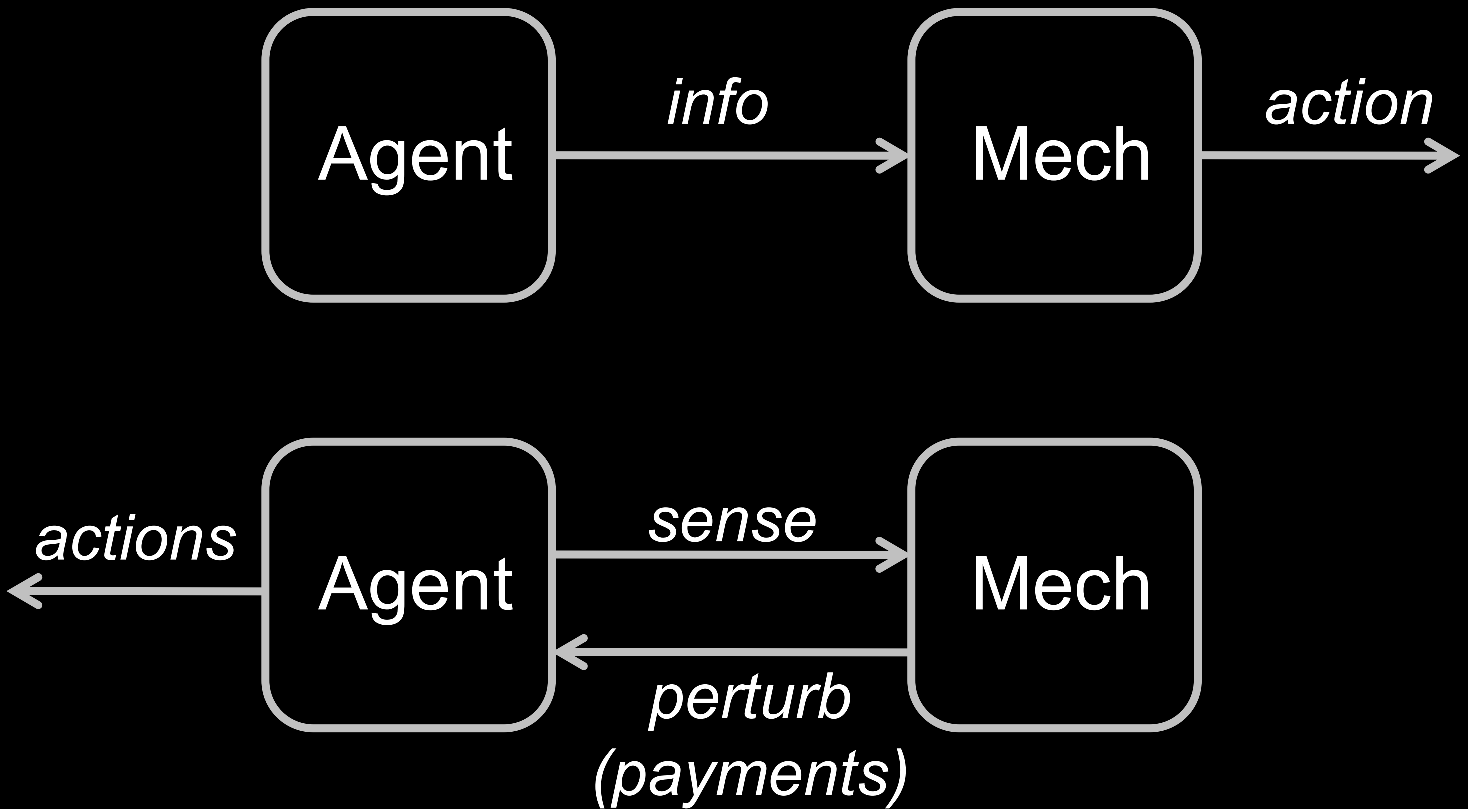

Environment Design

(Zhang & Parkes'08)

Role for AI

AI + crowdsourcing ≈ A New Kind of Firm

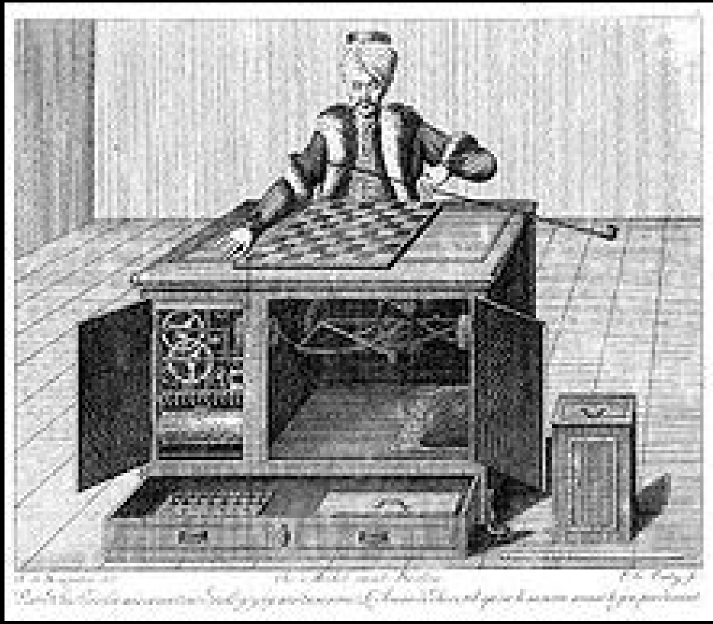

finally put the AI into the mechanical Turk?

Example: TopCoder

- Workers on TC get a *score* for a submission
  - correctness, docs, flexibility, extendability
  - combines to an aggregate “coder rating”

Example: TopCoder

(Archak'10)

- Workers on TC get a *score* for a submission
  - correctness, docs, flexibility, extendability
  - combines to an aggregate "coder rating"
- Skilled contestants tend to enter early
  - an implicit coordination mechanism
  - *signaling game*

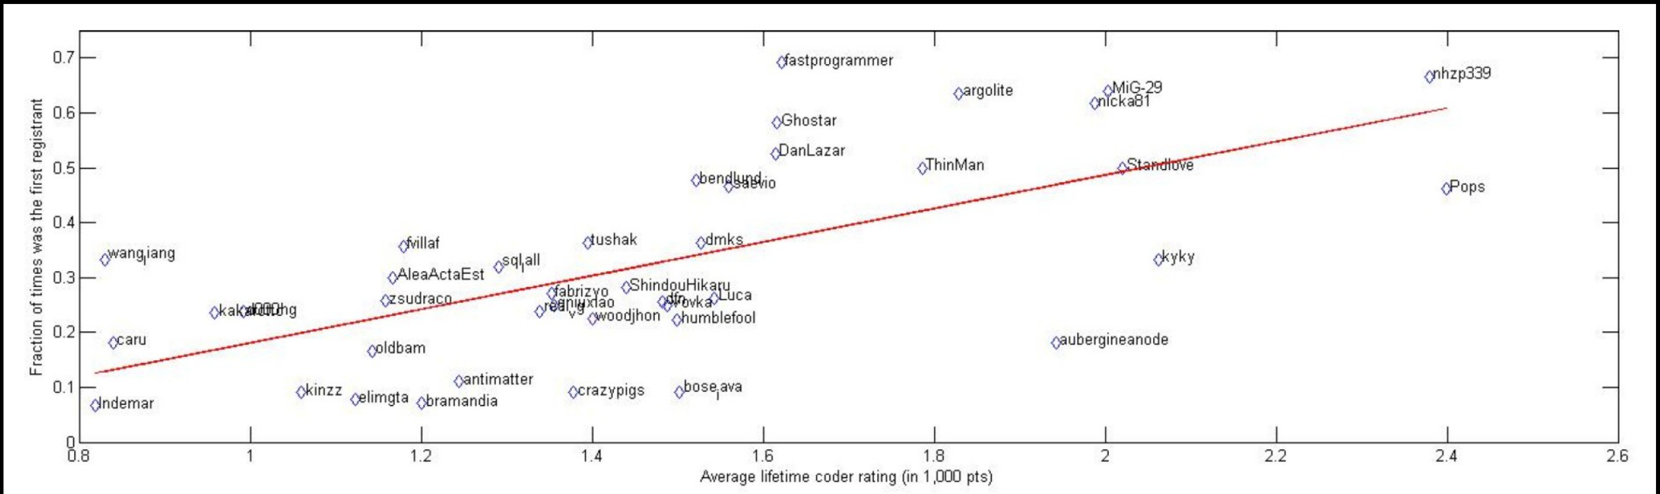

Generalized Task Markets

(Shahaf & Horvitz'10)

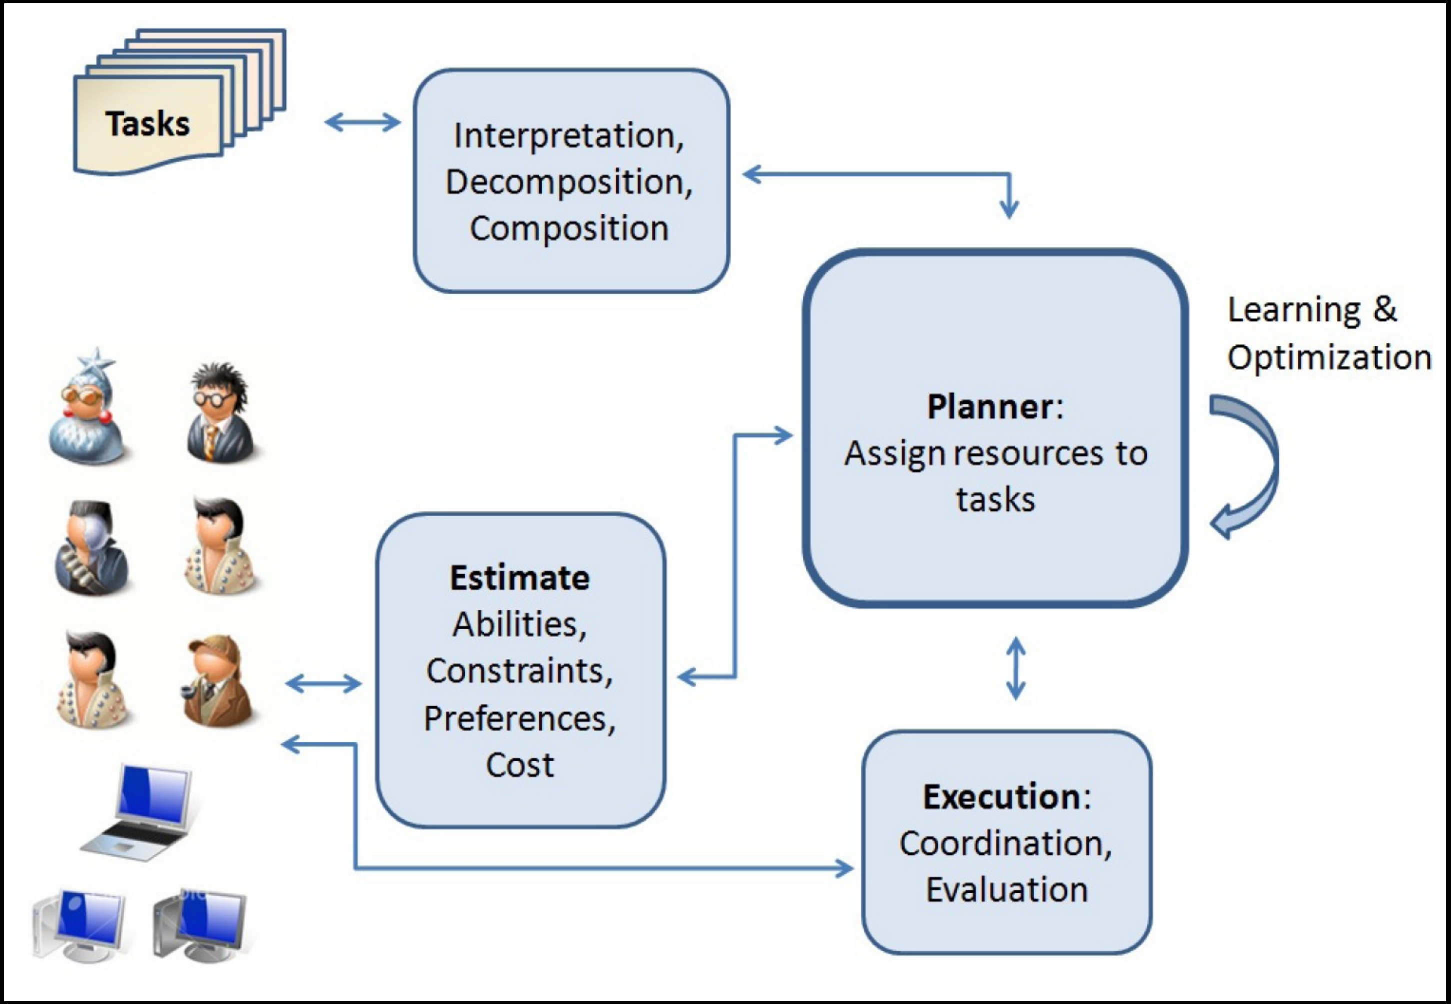

Example: Language Translation

(Shahaf & Horvitz'10)

- 388 participants, 70 countries, random trans. tasks
- Assign tasks to coalitions to maximize final quality while respecting capacity constraints

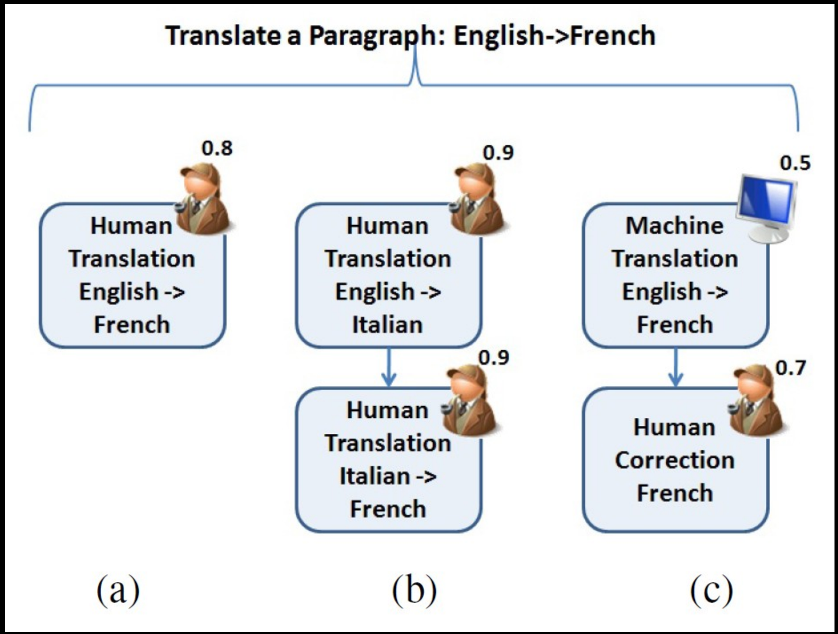

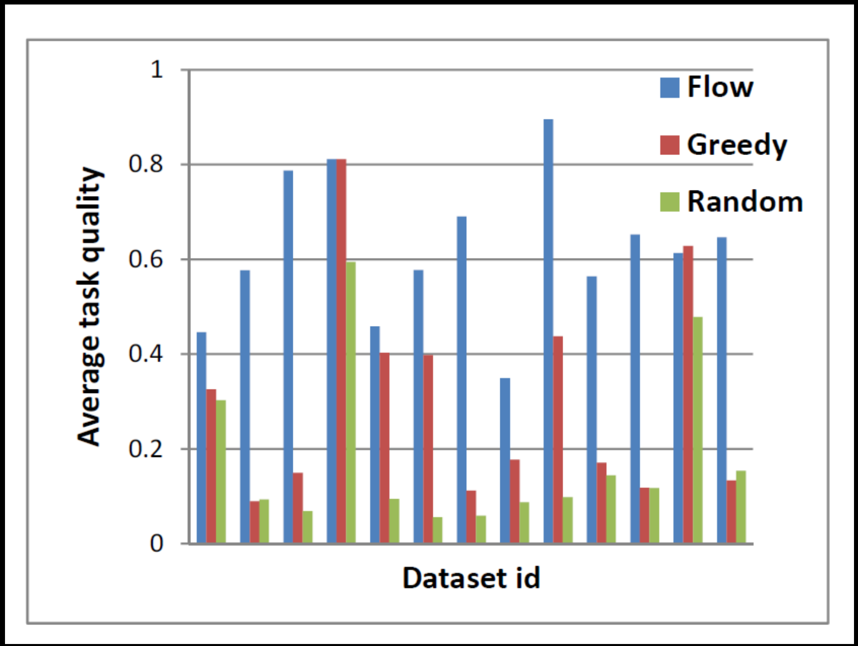

Example: Policy Teaching

(Zhang & Parkes'08, Zhang, Parkes & Chen'09)

MDP observe π perturb R $\rightarrow$ R+Δ Target policy π_T

Example: Policy Teaching

(Zhang & Parkes'08, Zhang, Parkes & Chen'09)

MDP   observe π perturb R $\rightarrow$ R+Δ Target policy π_T

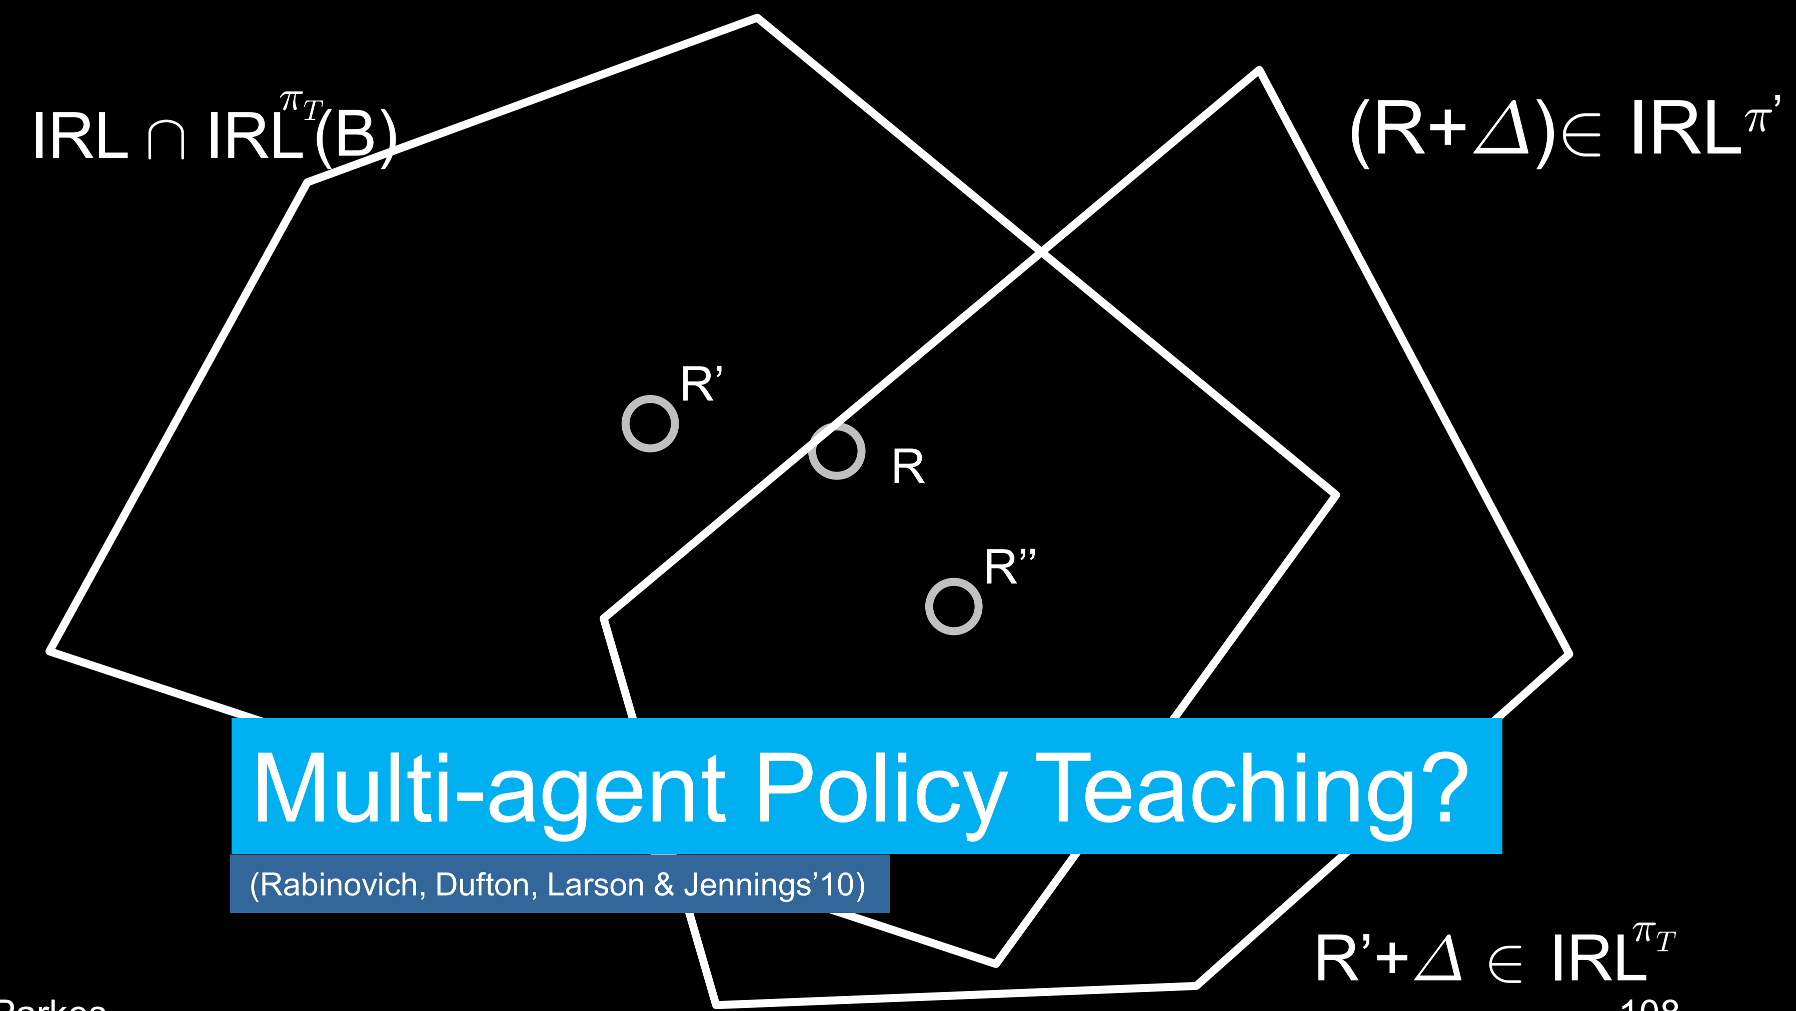

Multi-agent Policy Teaching?

(Rabinovich, Dufton, Larson & Jennings'10)

Wanted: Better User Models

(Horton & Chilton'10)

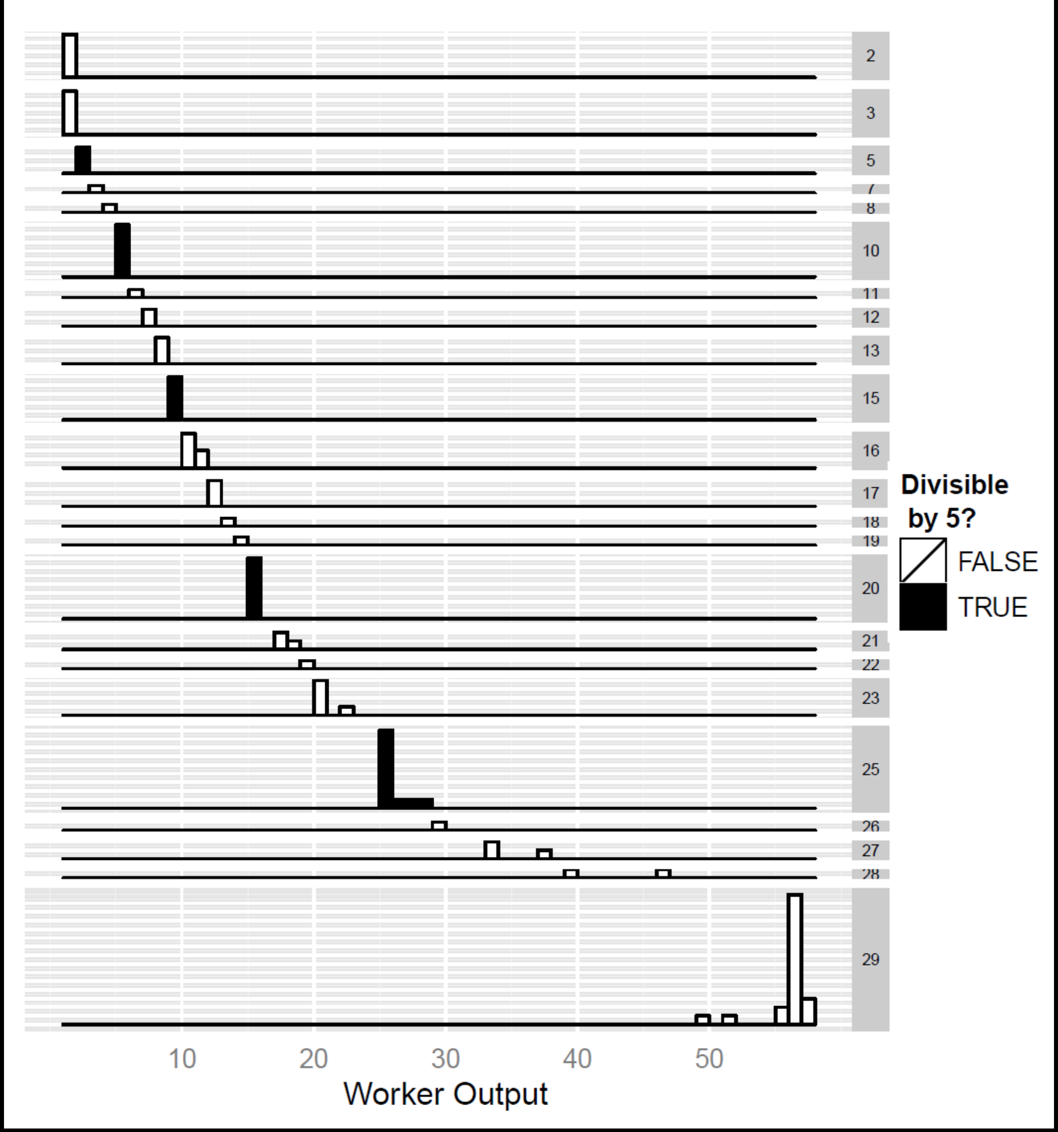

"target earnings" shows preference for amounts divisible by 5 cents

Hutong Karma

By midmorning, the vendors are out. They pedal through the alley on three-wheeled carts, each announcing his product with a trademark cry. The beer woman is the loudest, singing out again and again, "Maaaaiiiii piiiiijiuuuuuu!"... The rice man's refrain is higher-pitched; the vinegar dealer occupies the lower registers. ... The sounds are soothing, a reminder that even if I never left my doorway again life would be sustainable, albeit imbalanced. I would have cooking oil, soy sauce, and certain vegetables and fruit in season. In winter, I could buy strings of garlic. ...

On an average day, a recycler passes through every half hour, riding a flat-bed tricycle. ... Not long ago, I piled some useless possessions in the entryway of my apartment ... A stack of old magazines sold for sixty-two cents; a burned-out computer cord went for a nickel. Two broken lamps were seven cents, total. A worn-out pair of shoes: twelve cents. Two broken Palm Pilots: thirty-seven cents.

— *Hutong Karma. The many incarnations of a Beijing alleyway*, by Peter Hessler, The New Yorker, February 13, 2006.

Computational Sustainability through "Sharing Markets"

Sharing Markets

- Goal: use AI and electronic markets to transform our use of resources
- Support "microtransactions".
- For well functioning systems, need for:
  - scrip (Friedman, Halpern & Kash'06)
  - reputation (Friedman, Resnick & Sami'07)
  - accounting (Seuken, Tang & Parkes'10)

… and handle complexity!

Hidden Markets

(Seuken, Jain, Tan & Czerwinski ’10, Seuken, Jain & Parkes’10)

Hidden Markets

(Seuken, Jain, Tan & Czerwinski '10, Seuken, Jain & Parkes'10)

Example: P2P backup

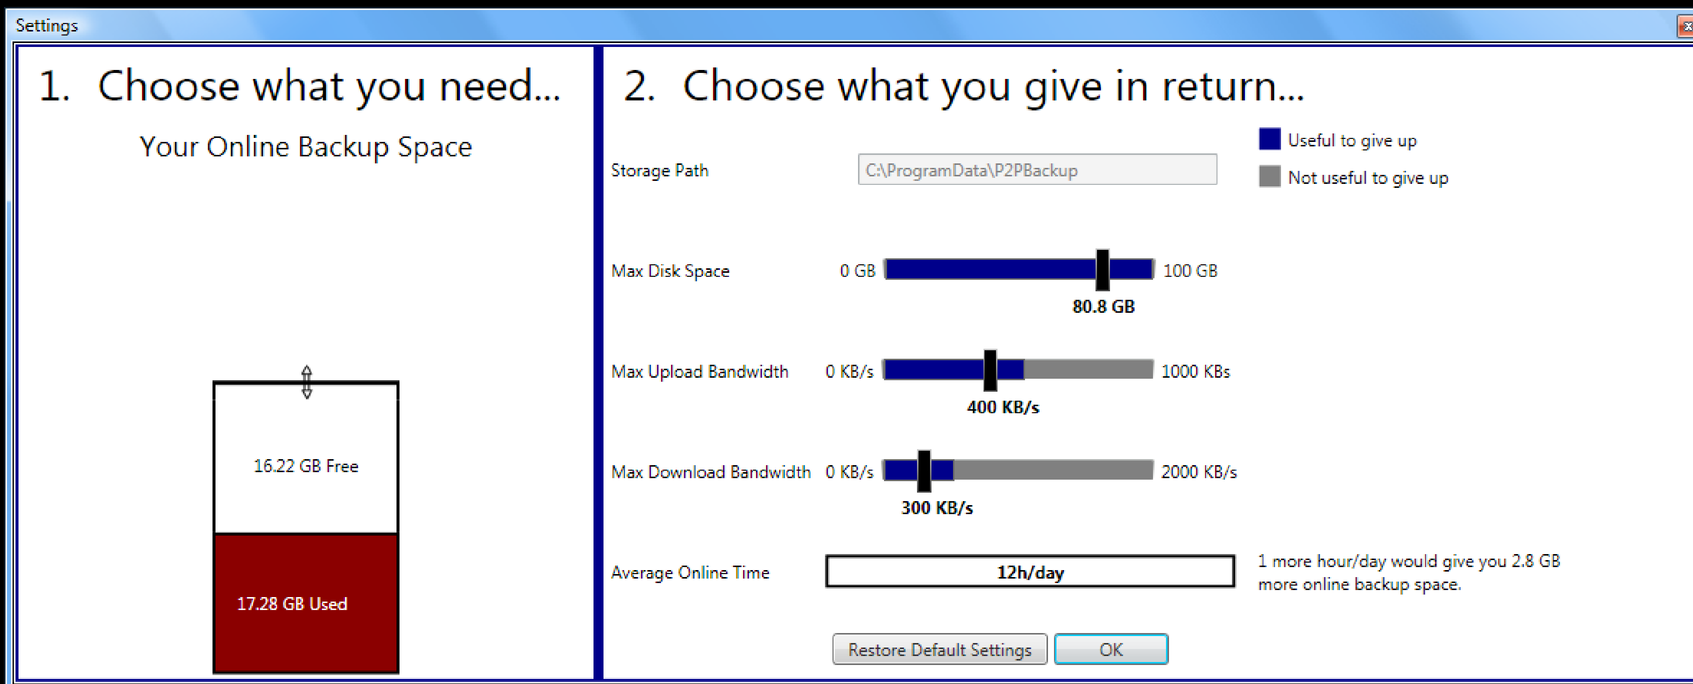

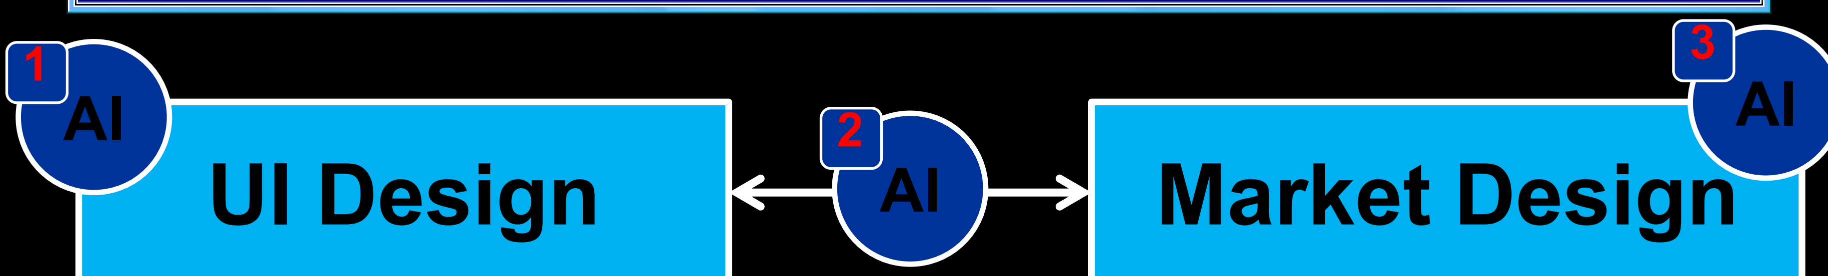

Summary

- MD theory is beautiful but severely stretched by Internet scale systems
- Provide useful formalism, but to make real progress in AI we'll need to move beyond
- Emphasized here three things:
  - heuristic approaches for MD
  - dynamic coordination opportunities
  - future: intelligent task and sharing markets

Acknowledgments

- People: David Bacon, Jonathan Bredin, Ruggiero Cavallo, David Chen, Yiling Chen, Florin Constantin, Jacomo Corbo, Quang Duong, Sujit Gujar, Boi Faltings, Felix Fischer, Eric Horvitz, Takayuki Ito, Shaili Jain, Adam Juda, Jayant Kalagnanam, Laura Kang, Ian Kash, Mark Klein, Jerry Kung, Sebastien Lahaie, John Lai, Ben Lubin, Debasis Mishra, Chaki Ng, Johan Pouwelse, Adrian Petcu, Ariel Procaccia, Malvika Rao, Tuomas Sandholm, Sven Seuken, Jeff Shneidman, Satinder Singh, Jie Tang, Chris Thorpe, Lyle Ungar, Dimah Yanovsky, Haoqi Zhang, James Zou.
- Grants: NSF, Microsoft, IBM, Yahoo!, ARL

www.eecs.harvard.edu/econcs

References

- (Moulin'80) On strategy-proofness and single peakedness, H. Moulin, Public Choice 35: 437-455 (1980)
  (Vickrey'61) Counterspeculation, Auctions, and Competitive Sealed Tenders. William Vickrey. Journal of Finance, 16:8-37 (1961)
  (Hurwicz'60) Optimality and informational efficiency in resource allocation processes, L. Hurwicz, in Arrow, Karlin and Suppes (eds.), Mathematical Methods in the Social Sciences. Stanford University Press (1960).
  (Hurwicz'72) On informationally decentralized systems, L. Hurwicz, in Radner and McGuire, Decision and Organization. North-Holland, Amsterdam.
  (Rosenschein & Zlotkin'94) Designing Conventions for Automated Negotiation, J. S. Rosenschein and G. Zlotkin, AI Magazine 15(3): 29-46 (1994)
  (Ephrati & Rosenschein'91) The Clarke Tax as a Consensus Mechanism Among Automated Agents, E. Ephrati, J. S. Rosenschein, Proc. 9th AAAI 1991: 173-178
  (Varian'95) Economic Mechanism Design for Computerized Agents, H. Varian, In USENIX workshop on Electronic Commerce (1995)
  (Goto '98) Goto.com
  (Edelman & Ostrovsky'07) Strategic bidder behavior in sponsored search auctions, B. Edelman and M. Ostrovsky, Decision Support Systems 43:192-198 (2007)
  (Rassenti, Smith and Bulfin'82) A combinatorial auction mechanism for airport time slot allocation, S. Rassenti, V. Smith and R. Bulfin, Bell Journal of Economics, 13, 402–417 (1982)
  (Sandholm'99) An Algorithm for Optimal Winner Determination in Combinatorial Auctions, T. Sandholm, Proc. 16th IJCAI'99: 542-547 (1999)
  (Fujishima et al.'99) Taming the Computational Complexity of Combinatorial Auctions: Optimal and Approximate Approaches, Y. Fujishima, K. Leyton-Brown and Y. Shoham, Proc. 16th IJCAI'99: 548-553 (1999)
  (Nisan'00) Bidding and allocation in combinatorial auctions, N. Nisan, Proc. 2nd ACM Conf. on E Commerce (EC'00): 1-12 (2000)
  (Boutilier & Hoos'01) Bidding Languages for Combinatorial Auctions, C. Boutilier and H. Hoos, Proc. 17th IJCAI'01:1211-1217 (2001)

References

- (Rothkopf et al. '98) Computationally Manageable Combinational Auctions
M. H. Rothkopf, A. Pekec, R. M. Harstad, Management Science 44: 1131-1147 (1998)
(Parkes & Ungar'00) Iterative Combinatorial Auctions: Theory and Practice, David C. Parkes and Lyle H. Ungar, Proc. 17th AAAI'00: 74-81 (2000)
(Lahaie & Parkes'04) Applying Learning Algorithms to Preference Elicitation, S. Lahaie and D. C. Parkes, Proc. 5th ACM Conference on Electronic Commerce (EC'04): 180-188, 2004
(Hudson & Sandholm'03) Using value queries in combinatorial auctions, B. Hudson and T. Sandholm, Proc. 4th ACM Conference on Electronic Commerce (EC'03), 2003
(Hyafil & Boutilier'06) Regret-based incremental partial revelation mechanisms, N. Hyafil and C. Boutilier. In Proc. AAAI-06: 672-678, 2006.
(Nisan & Ronen'99) Algorithmic Mechanism Design, N. Nisan and A. Ronen, Proc. 31st STOC 1999: 129-140 (1999).
(Lehmann et al.'99) Truth revelation in approximately efficient combinatorial auctions, D. Lehmann, L. I. O'Callaghan, and Y. Shoham Proc. 1st ACM conference on Electronic commerce (EC'99):96-102, 1999
(Likhodedov and Sandholm'04) Methods for Boosting Revenue in Combinatorial Auctions, Likhodedov, A. and Sandholm, T., In Proc. AAAI'04 (2004)
(Guo & Conitzer'08) Optimal-in-Expectation Redistribution Mechanisms, M. Guo and V. Conitzer, Proc. 7th AAMAS'08, 2008
(van Hentenryck and Bent'06) Online Stochastic Combinatorial Optimization. P. Van Hentenryck and R.Bent. MIT Press, 2006.
(Parkes & Duong'07) An Ironing-Based Approach to Adaptive Online Mechanism Design in Single-Valued Domains, D. C. Parkes and Q. Duong, in Proc. 22nd National Conference on Artificial Intelligence (AAAI'07), 94-101, 2007
(Constantin & Parkes'09) Self-Correcting Sampling-Based Dynamic Multi-Unit Auctions, F. Constantin and D. C. Parkes, in Proc. 10th ACM Electronic Commerce Conference (EC'09), pages 89-98, 2009
(Roth'03) The Origins, History, and Design of the Resident Match, A E. Roth, Journal of the American Medical Association, 289: 909-212 (2003)
(Pathak and Sonmez'08) Leveling the Playing Field: Sincere and Strategic Players in the Boston Mechanism, P. Pathak and T. Sonmez, American Economic Review (2008).
(Rastegeri, Condon & Leyton-Brown'10) Revenue Monotonicity in Deterministic, Dominant-Strategy Combinatorial Auctions, B. Rastegeri, A. Condon, K. Leyton-Brown, Artificial Intelligence (AIJ), 2010

References

- (Pathak and Sonmez'09) Comparing Mechanisms by their Vulnerability to Manipulation, P. A. Pathak and T. Sonmez, 2009
  (Ausubel and Milgrom'06) The Lovely but Lonely Vickrey Auction, L. M. Ausubel and P. Milgrom, in Nisan et al. Combinatorial Auctions (MIT Press) 2006
  (Parkes, Kalagnanam and Esi'01) Achieving Budget-Balance with Vickrey-Based Payment Schemes in Exchanges, D. C. Parkes, J. R. Kalagnanam and M. Eso,in Proc. 17th International Joint Conference on Artificial Intelligence (IJCAI'01), pages 1161-1168, 2001.
  (Lubin & Parkes'09) Quantifying the Strategyproofness of Mechanisms via Metrics on Payoff Distributions, B. Lubin and D. C. Parkes, Proc. 25th UAI'09: 349-358
  (Parkes & Singh'03) An MDP-Based Approach to Online Mechanism Design, D. C. Parkes and S. Singh, in Proc. 17th Annual Conf. on Neural Information Processing Systems (NIPS'03), 2003
  (Parkes, Singh & Yanovsky'04) Approximately Efficient Online Mechanism Design, D. C. Parkes, S. Singh, and D. Yanovsky, in Proc. 18th Annual Conf. on Neural Information Processing Systems (NIPS'04), 2004
  (Bergemann & Valimaki'08) The dynamic pivot mechanism, D. Bergemann and J. Valimaki. Cowles Foundation Discussion Paper 1672, 2008
  (Cavallo, Parkes & Singh'09) Efficient Mechanisms with Dynamic Populations and Dynamic Types, R. Cavallo, D. C. Parkes, and S. Singh, 2009
  (Cavallo & Parkes'08) Efficient Metadeliberation Auctions, R. Cavallo and D. C. Parkes, Proc. 23rd AAAI Conference on Artificial Intelligence (AAAI'08), Chicago, IL, pages 50-56, 2008
  (Gerding, Stein, Larson, Rogers and Jennings'10) Scaleable mechanism design for the procurement of services with uncertain durations, E. Gerding, S. Stein, K. Larson, A. Rogers and N. R. Jennings, Proc. 9th Int Conf. on Autonomous Agents and Multi-Agent Systems, (2010)
  (Zhang & Parkes'08) Value-based Policy Teaching with Active Indirect Elicitation, H. Zhang and D. C. Parkes, Proc. 23rd AAAI Conference on Artificial Intelligence (AAAI'08), 208-214, 2008
  (Archak'10) Money, glory and cheap talk: analyzing strategic behavior of contestants in simultaneous crowdsourcing contests on TopCoder.com, N. Archak, Proc. 19th International World Wide Web Conference (WWW 2010), 2010
  (Shahaf & Horvitz'10) Generalized Task Markets for Human and Machine Computation, D. Shahaf and E. Horvitz, Proc. AAAI'10 (2010)

References

- (Zhang, Parkes & Chen'09) Policy Teaching Through Reward Function Learning, H. Zhang, D. C. Parkes and Y. Chen, Proc. 10th ACM Electronic Commerce Conference (EC'09), pages 295-304, 2009
  (Horton & Chilton'10) The Labor Economics of Paid Crowdsourcing, J. Horton and L. Chilton, Proc. 11th ACM Conf. on E. Commerce (EC'10), 2010
  (Friedman, Halpern & Kash'06) Efficiency and Nash Equilibria in a Scrip System for P2P Networks, E. Friedman, J. Halpern and I. Kash, Proc. 7th ACM Conference on Electronic Commerce (EC'06), pages 140-149, 2006
  (Friedman, Resnick & Sami'07) Manipulation-Resistant Reputation Systems, E. Friedman, P. Resnick, R. Sami, in Algorithmic Game Theory, (N. Nisan, T. Roughgarden, E. Tardos, V. Vazirani, editors), Cambridge University Press, 2007.
  (Seuken, Tang & Parkes'10) Accounting Mechanisms for Distributed Work Systems, S. Seuken, J. Tang, and D. C. Parkes, Proc. 24th AAAI Conference on Artificial Intelligence (AAAI'10), 2010
  (Seuken, Jain, Tan & Czerwinski'10) Hidden Markets: UI Design for a P2P Backup Application, S. Seuken, K. Jain, D. Tan and M. Czerwinski, In Proc. Conference on Human Factors in Computing Systems (CHI'10), 2010
  (Seuken, Jain & Parkes'10) Hidden Market Design, S. Seuken, K. Jain and D. C. Parkes, Proc. 24th AAAI Conference on Artificial Intelligence (AAAI'10), 2010